

# Hypernyms under Siege: Linguistically-motivated Artillery for Hypernymy Detection

Vered Shwartz<sup>1</sup>, Enrico Santus<sup>2,3</sup> and Dominik Schlechtweg<sup>4</sup>

<sup>1</sup>Bar-Ilan University, Ramat-Gan, Israel

<sup>2</sup>Singapore University of Technology and Design, Singapore

<sup>3</sup>The Hong Kong Polytechnic University, Hong Kong

<sup>4</sup>University of Stuttgart, Stuttgart, Germany

{vered1986, esantus}@gmail.com, dominik.schlechtweg@gmx.de

## Abstract

The fundamental role of hypernymy in NLP has motivated the development of many methods for the automatic identification of this relation, most of which rely on word distribution. We investigate an extensive number of such unsupervised measures, using several distributional semantic models that differ by context type and feature weighting. We analyze the performance of the different methods based on their linguistic motivation. Comparison to the state-of-the-art supervised methods shows that while supervised methods generally outperform the unsupervised ones, the former are sensitive to the distribution of training instances, hurting their reliability. Being based on general linguistic hypotheses and independent from training data, unsupervised measures are more robust, and therefore are still useful artillery for hypernymy detection.

## 1 Introduction

In the last two decades, the NLP community has invested a consistent effort in developing automated methods to recognize hypernymy. Such effort is motivated by the role this semantic relation plays in a large number of tasks, such as taxonomy creation (Snow et al., 2006; Navigli et al., 2011) and recognizing textual entailment (Dagan et al., 2013). The task has appeared to be, however, a challenging one, and the numerous approaches proposed to tackle it have often shown limitations.

Early corpus-based methods have exploited patterns that may indicate hypernymy (e.g. “*animals* such as *dogs*”) (Hearst, 1992; Snow et al., 2005), but the recall limitation of this approach, requiring both words to co-occur in a sentence, motivated the development of methods that rely on

adaptations of the *distributional hypothesis* (Harris, 1954).

The first distributional approaches were unsupervised, assigning a score for each  $(x, y)$  word-pair, which is expected to be higher for hypernym pairs than for negative instances. Evaluation is performed using ranking metrics inherited from information retrieval, such as Average Precision (AP) and Mean Average Precision (MAP). Each measure exploits a certain linguistic hypothesis such as the *distributional inclusion hypothesis* (Weeds and Weir, 2003; Kotlerman et al., 2010) and the *distributional informativeness hypothesis* (Santus et al., 2014; Rimell, 2014).

In the last couple of years, the focus of the research community shifted to supervised distributional methods, in which each  $(x, y)$  word-pair is represented by a combination of  $x$  and  $y$ 's word vectors (e.g. concatenation or difference), and a classifier is trained on these resulting vectors to predict hypernymy (Baroni et al., 2012; Roller et al., 2014; Weeds et al., 2014). While the original methods were based on count-based vectors, in recent years they have been used with word embeddings (Mikolov et al., 2013; Pennington et al., 2014), and have gained popularity thanks to their ease of use and their high performance on several common datasets. However, there have been doubts on whether they can actually learn to recognize hypernymy (Levy et al., 2015b).

Additional recent hypernymy detection methods include a multimodal perspective (Kiela et al., 2015), a supervised method using unsupervised measure scores as features (Santus et al., 2016a), and a neural method integrating path-based and distributional information (Shwartz et al., 2016).

In this paper we perform an extensive evaluation of various unsupervised distributional measures for hypernymy detection, using several distributional semantic models that differ by context type and feature weighting. Some measure vari-

ants and context-types are tested for the first time.<sup>1</sup>

We demonstrate that since each of these measures captures a different aspect of the hypernymy relation, there is no single measure that consistently performs well in discriminating hypernymy from different semantic relations. We analyze the performance of the measures in different settings and suggest a principled way to select the suitable measure, context type and feature weighting according to the task setting, yielding consistent performance across datasets.

We also compare the unsupervised measures to the state-of-the-art supervised methods. We show that supervised methods outperform the unsupervised ones, while also being more efficient, computed on top of low-dimensional vectors. At the same time, however, our analysis reassesses previous findings suggesting that supervised methods do not actually learn the relation between the words, but only characteristics of a single word in the pair (Levy et al., 2015b). Moreover, since the features in embedding-based classifiers are latent, it is difficult to tell what the classifier has learned. We demonstrate that unsupervised methods, on the other hand, do account for the relation between words in a pair, and are easily interpretable, being based on general linguistic hypotheses.


## 2 Distributional Semantic Spaces

We created multiple distributional semantic spaces that differ in their context type and feature weighting. As an underlying corpus we used a concatenation of the following two corpora: `ukWaC` (Ferraresi, 2007), a 2-billion word corpus constructed by crawling the `.uk` domain, and `WaCkypedia_EN` (Baroni et al., 2009), a 2009 dump of the English Wikipedia. Both corpora include POS, lemma and dependency parse annotations. Our vocabulary (of target and context words) includes only nouns, verbs and adjectives that occurred at least 100 times in the corpus.

**Context Type** We use several context types:

- **Window-based contexts:** the contexts of a target word  $w_i$  are the words surrounding it in a  $k$ -sized window:  $w_{i-k}, \dots, w_{i-1}, w_{i+1}, \dots, w_{i+k}$ . If the context-type is directional, words occurring before and after  $w_i$  are marked differently, i.e.:  $w_{i-k}/l, \dots, w_{i-1}/l, w_{i+1}/r, \dots, w_{i+k}/r$ .

<sup>1</sup>Our code and data are available at:  
<https://github.com/vered1986/UnsupervisedHypernymy>



**Figure 1:** An example dependency tree of the sentence *cute cats drink milk*, with the target word *cats*. The dependency-based contexts are *drink-v:nsbj* and *cute-a:amod*<sup>-1</sup>. The joint-dependency context is *drink-v#milk-n*. Differently from Chersoni et al. (2016), we exclude the dependency tags to mitigate the sparsity of contexts.

Out-of-vocabulary words are filtered out before applying the window. We experimented with window sizes 2 and 5, directional and indirectional (`win2`, `win2d`, `win5`, `win5d`).

- **Dependency-based contexts:** rather than adjacent words in a window, we consider neighbors in a dependency parse tree (Padó and Lapata, 2007; Baroni and Lenci, 2010). The contexts of a target word  $w_i$  are its parent and daughter nodes in the dependency tree (`dep`). We also experimented with a joint dependency context inspired by Chersoni et al. (2016), in which the contexts of a target word are the parent-sister pairs in the dependency tree (`joint`). See Figure 1 for an illustration.

**Feature Weighting** Each distributional semantic space is spanned by a matrix  $M$  in which each row corresponds to a target word while each column corresponds to a context. The value of each cell  $M_{i,j}$  represents the association between the target word  $w_i$  and the context  $c_j$ . We experimented with two feature weightings:

- **Frequency** - raw frequency (no weighting):  $M_{i,j}$  is the number of co-occurrences of  $w_i$  and  $c_j$  in the corpus.
- **Positive PMI (PPMI)** - pointwise mutual information (PMI) (Church and Hanks, 1990) is defined as the log ratio between the joint probability of  $w$  and  $c$  and the product of their marginal probabilities:  $PMI(w, c) = \log \frac{\hat{P}(w, c)}{\hat{P}(w)\hat{P}(c)}$ , where  $\hat{P}(w)$ ,  $\hat{P}(c)$ , and  $\hat{P}(w, c)$  are estimated by the relative frequencies of a word  $w$ , a context  $c$  and a word-context pair  $(w, c)$ , respectively. To handle unseen pairs  $(w, c)$ , yielding  $PMI(w, c) = \log(0) = -\infty$ , PPMI (Bullinaria and Levy, 2007) assigns zero to negative PMI scores:  $PPMI(w, c) = \max(PMI(w, c), 0)$ .

In addition, one of the measures we used (San-tus et al., 2014) required a third feature weighting:

- **Positive LMI (PLMI)** - positive local mutual information (PLMI) (Evert, 2005; Evert, 2008). PPMI was found to have a bias towards rare events. PLMI simply balances PPMI by multiplying it by the co-occurrence frequency of  $w$  and  $c$ :  $PLMI(w, c) = freq(w, c) \cdot PPMI(w, c)$ .

### 3 Unsupervised Hypernymy Detection Measures

We experiment with a large number of unsupervised measures proposed in the literature for distributional hypernymy detection, with some new variants. In the following section,  $\vec{v}_x$  and  $\vec{v}_y$  denote  $x$  and  $y$ 's word vectors (rows in the matrix  $M$ ). We consider the scores as measuring to what extent  $y$  is a hypernym of  $x$  ( $x \rightarrow y$ ).

#### 3.1 Similarity Measures

Following the *distributional hypothesis* (Harris, 1954), similar words share many contexts, thus have a high similarity score. Although the hypernymy relation is asymmetric, similarity is one of its properties (Santus et al., 2014).

- **Cosine Similarity** (Salton and McGill, 1986) A symmetric similarity measure:

$$\cos(x, y) = \frac{\vec{v}_x \cdot \vec{v}_y}{\|\vec{v}_x\| \cdot \|\vec{v}_y\|}$$

- **Lin Similarity** (Lin, 1998) A symmetric similarity measure that quantifies the ratio of shared contexts to the contexts of each word:

$$Lin(x, y) = \frac{\sum_{c \in \vec{v}_x \cap \vec{v}_y} [\vec{v}_x[c] + \vec{v}_y[c]]}{\sum_{c \in \vec{v}_x} \vec{v}_x[c] + \sum_{c \in \vec{v}_y} \vec{v}_y[c]}$$

- **APSyn** (Santus et al., 2016b) A symmetric measure that computes the extent of intersection among the  $N$  most related contexts of two words, weighted according to the rank of the shared contexts (with  $N$  as a hyper-parameter):

$$APSyn(x, y) = \sum_{c \in N(\vec{v}_x) \cap N(\vec{v}_y)} \frac{1}{\frac{rank_x(c) + rank_y(c)}{2}}$$

#### 3.2 Inclusion Measures

According to the *distributional inclusion hypothesis*, the prominent contexts of a hyponym ( $x$ ) are expected to be included in those of its hypernym ( $y$ ).

- **Weeds Precision** (Weeds and Weir, 2003) A directional precision-based similarity measure. This measure quantifies the weighted inclusion of  $x$ 's contexts by  $y$ 's contexts:

$$WeedsPrec(x \rightarrow y) = \frac{\sum_{c \in \vec{v}_x \cap \vec{v}_y} \vec{v}_x[c]}{\sum_{c \in \vec{v}_x} \vec{v}_x[c]}$$

- **cosWeeds** (Lenci and Benotto, 2012) Geometric mean of cosine similarity and Weeds precision:

$$\cosWeeds(x \rightarrow y) = \sqrt{\cos(x, y) \cdot WeedsPrec(x \rightarrow y)}$$

- **ClarkeDE** (Clarke, 2009) Computes degree of inclusion, by quantifying weighted coverage of the hyponym's contexts by those of the hypernym:

$$CDE(x \rightarrow y) = \frac{\sum_{c \in \vec{v}_x \cap \vec{v}_y} \min(\vec{v}_x[c], \vec{v}_y[c])}{\sum_{c \in \vec{v}_x} \vec{v}_x[c]}$$

- **balAPinc** (Kotlerman et al., 2010) Balanced average precision inclusion.

$$APinc(x \rightarrow y) = \frac{\sum_{r=1}^{N_y} [P(r) \cdot rel(c_r)]}{N_y}$$

is an adaptation of the average precision measure from information retrieval for the inclusion hypothesis.  $N_y$  is the number of non-zero contexts of  $y$  and  $P(r)$  is the precision at rank  $r$ , defined as the ratio of shared contexts with  $y$  among the top  $r$  contexts of  $x$ .  $rel(c)$  is the relevance of a context  $c$ , set to 0 if  $c$  is not a context of  $y$ , and to  $1 - \frac{rank_y(c)}{N_y + 1}$  otherwise, where  $rank_y(c)$  is the rank of the context  $c$  in  $y$ 's sorted vector. Finally,

$$balAPinc(x \rightarrow y) = \sqrt{Lin(x, y) \cdot APinc(x \rightarrow y)}$$

is the geometric mean of APinc and Lin similarity.

- **invCL** (Lenci and Benotto, 2012) Measures both distributional inclusion of  $x$  in  $y$  and distributional non-inclusion of  $y$  in  $x$ :

$$invCL(x \rightarrow y) = \sqrt{CDE(x \rightarrow y) \cdot (1 - CDE(y \rightarrow x))}$$

#### 3.3 Informativeness Measures

According to the *distributional informativeness hypothesis*, hypernyms tend to be less informative than hyponyms, as they are likely to occur in more general contexts than their hyponyms.

- **SLQS** (Santus et al., 2014)

$$SLQS(x \rightarrow y) = 1 - \frac{E_x}{E_y}$$

The informativeness of a word  $x$  is evaluated as the median entropy of its top  $N$  contexts:  $E_x = \text{median}_{i=1}^N (H(c_i))$ , where  $H(c)$  is the entropy of context  $c$ .

- **SLQS Sub** A new variant of SLQS based on the assumption that if  $y$  is judged to be a hypernym of  $x$  to a certain extent, then  $x$  should be judged to be a hyponym of  $y$  to the same extent (which is not the case for regular SLQS). This is achieved by subtraction:

$$SLQS_{sub}(x \rightarrow y) = E_y - E_x$$

It is weakly symmetric in the sense that  $SLQS_{sub}(x \rightarrow y) = -SLQS_{sub}(y \rightarrow x)$ .

SLQS and SLQS Sub have 3 hyper-parameters: i) the number of contexts  $N$ ; ii) whether to use median or average entropy among the top  $N$  contexts; and iii) the feature weighting used to sort the contexts by relevance (i.e., PPMI or PLMI).

- **SLQS Row** Differently from SLQS, SLQS Row computes the entropy of the target rather than the average/median entropy of the contexts, as an alternative way to compute the generality of a word.<sup>2</sup> In addition, parallel to SLQS we tested SLQS Row with subtraction, **SLQS Row Sub**.
- **RCTC** (Rimell, 2014) Ratio of change in topic coherence:

$$RCTC(x \rightarrow y) = \frac{TC(t_x)/TC(t_{x \setminus y})}{TC(t_y)/TC(t_{y \setminus x})}$$

where  $t_x$  are the top  $N$  contexts of  $x$ , considered as  $x$ 's *topic*, and  $t_{x \setminus y}$  are the top  $N$  contexts of  $x$  which are not contexts of  $y$ .  $TC(A)$  is the topic coherence of a set of words  $A$ , defined as the median pairwise PMI scores between words in  $A$ .  $N$  is a hyper-parameter. The measure is based on the assumptions that excluding  $y$ 's contexts from  $x$ 's increases the coherence of the topic, while excluding  $x$ 's contexts from  $y$ 's decreases the coherence of the topic. We include this measure under the informativeness inclusion, as it is based on a similar hypothesis.

### 3.4 Reversed Inclusion Measures

These measures are motivated by the fact that, even though—being more general—hypernyms are expected to occur in a larger set of contexts, sentences like “the *vertebrate* barks” or “the *mammal* arrested the thieves” are not common, since hyponyms are more specialized and are hence more appropriate in such contexts. On the other

<sup>2</sup>In our preliminary experiments, we noticed that the entropies of the targets and those of the contexts are not highly correlated, yielding a Spearman’s correlation of up to 0.448 for window based spaces, and up to 0.097 for the dependency-based ones ( $p < 0.01$ ).

dataset	relations	#instances	size
BLESS	hypernym	1,337	26,554
	meronym	2,943	
	coordination	3,565	
	event	3,824	
	attribute	2,731	
	random-n	6,702	
	random-v	3,265	
EVALution	hypernym	2,187	13,465 <sup>3</sup>
	hypernym	3,637	
	meronym	1,819	
	attribute	2,965	
	synonym	1,888	
Lenci/Benotto	antonym	3,156	5,010
	hypernym	1,933	
	synonym	1,311	
Weeds	antonym	1,766	2,928
	hypernym	1,469	
	coordination	1,459	

**Table 1:** The semantic relations, number of instances in each relation, and size of each dataset.

hand, hyponyms are likely to occur in broad contexts (e.g. *eat*, *live*), where hypernyms are also appropriate. In this sense, we can define the *reversed inclusion hypothesis*: “hypernym’s contexts are likely to be included in the hyponym’s contexts”. The following variants are tested for the first time.

- **Reversed Weeds**

$$RevWeeds(x \rightarrow y) = Weeds(y \rightarrow x)$$

- **Reversed ClarkeDE**

$$RevCDE(x \rightarrow y) = CDE(y \rightarrow x)$$

## 4 Datasets

We use four common semantic relation datasets: BLESS (Baroni and Lenci, 2011), EVALution (Santus et al., 2015), Lenci/Benotto (Benotto, 2015), and Weeds (Weeds et al., 2014). The datasets were constructed either using knowledge resources (e.g. WordNet, Wikipedia), crowd-sourcing or both. The semantic relations and the size of each dataset are detailed in Table 1.

In our distributional semantic spaces, a target word is represented by the word and its POS tag. While BLESS and Lenci/Benotto contain this information, we needed to add POS tags to the other datasets. For each pair  $(x, y)$ , we considered 3 pairs  $(x-p, y-p)$  for  $p \in \{noun, adjective, verb\}$ , and added the respective pair to the dataset only if the words were present in the corpus.<sup>4</sup>

<sup>3</sup>We removed the *entailment* relation, which had too few instances, and conflated relations to coarse-grained relations (e.g. *HasProperty* and *HasA* into *attribute*).

<sup>4</sup>Lenci/Benotto includes pairs to which more than one relation is assigned, e.g. when  $x$  or  $y$  are polysemous, and re-

dataset	hyper vs. relation	measure	context type	feature weighting	hyper-parameters	AP@100	AP@All
EVALution	all other relations	invCL	joint	freq	-	0.661	0.353
	meronym	APSyn	joint	freq	$N=500$	0.883	0.675
	attribute	APSyn	joint	freq	$N=500$	0.88	0.651
	antonym	SLQS_row	joint	freq	-	0.74	0.54
			joint	ppmi		0.74	0.55
			joint	plmi		0.74	0.537
	synonym	SLQS_row	joint	freq	-	0.83	0.647
			joint	ppmi		0.83	0.657
			joint	plmi		0.83	0.645
BLESS	all other relations	invCL	win5	freq	-	0.54	0.051
	meronym	SLQS <sub>sub</sub>	win5d	freq	$N=100, median, plmi$	1.0	0.76
		SLQS	win5d	freq	$N=100, median, plmi$	1.0	0.758
	coord	SLQS <sub>sub</sub>	joint	freq	$N=50, average, plmi$	0.995	0.537
	attribute	SLQS <sub>sub</sub>	dep	plmi	$N=70, average, plmi$	1.0	0.74
		cosine	joint	freq	-	1.0	0.622
	event	APSyn	dep	freq	$N=1000$	1.0	0.779
Lenci/ Benotto	all other relations	APSyn	joint	freq	$N=1000$	0.617	0.382
	antonym	APSyn	dep	freq	$N=1000$	0.861	0.624
	synonym	SLQS_row <sub>sub</sub>	joint	ppmi	-	0.948	0.725
Weeds	all other relations	clarkeDE	win5d	freq	-	0.911	0.441
	coord	clarkeDE	win5d	freq	-	0.911	0.441

**Table 2:** Best performing unsupervised measures on each dataset in terms of Average Precision (AP) at  $k = 100$ , for hypernym vs. all other relations and vs. each single relation. AP for  $k = all$  is also reported for completeness. We excluded the experiments of hypernym vs. random-(n, v, j) for brevity; most of the similarity and some of the inclusion measures achieve  $AP@100 = 1.0$  in these experiments.

We split each dataset randomly to 90% test and 10% validation. The validation sets are used to tune the hyper-parameters of several measures: SLQS (Sub), APSyn and RCTC.

## 5 Experiments

### 5.1 Comparing Unsupervised Measures

In order to evaluate the unsupervised measures described in Section 3, we compute the measure scores for each  $(x, y)$  pair in each dataset. We first measure the method’s ability to discriminate hypernymy from all other relations in the dataset, i.e. by considering hypernyms as positive instances, and other word pairs as negative instances. In addition, we measure the method’s ability to discriminate hypernymy from every other relation in the dataset by considering *one* relation at a time. For a relation  $\mathcal{R}$  we consider only  $(x, y)$  pairs that are annotated as either hypernyms (positive instances) or  $\mathcal{R}$  (negative instances). We rank the pairs according to the measure score and compute average precision (AP) at  $k = 100$  and  $k = all$ .<sup>5</sup>

lated differently in each sense. We consider  $y$  as a hypernym of  $x$  if hypernymy holds in some of the words’ senses. Therefore, when a pair is assigned both hypernymy and another relation, we only keep it as hypernymy.

<sup>5</sup>We tried several cut-offs and chose the one that seemed to be more informative in distinguishing between the unsupervised measures.

Table 2 reports the best performing measure(s), with respect to  $AP@100$ , for each relation in each dataset. The first observation is that there is no single combination of measure, context type and feature weighting that performs best in discriminating hypernymy from all other relations. In order to better understand the results, we focus on the second type of evaluation, in which we discriminate hypernyms from each other relation.

The results show preference to the syntactic context-types (`dep` and `joint`), which might be explained by the fact that these contexts are richer (as they contain both proximity and syntactic information) and therefore more discriminative. In feature weighting there is no consistency, but interestingly, raw frequency appears to be successful in hypernymy detection, contrary to previously reported results for word similarity tasks, where PPMI was shown to outperform it (Bullinaria and Levy, 2007; Levy et al., 2015a).

The new SLQS variants are on top of the list in many settings. In particular they perform well in discriminating hypernyms from symmetric relations (antonymy, synonymy, coordination).

The measures based on the *reversed inclusion hypothesis* performed inconsistently, achieving perfect score in the discrimination of hypernyms from unrelated words, and performing well

relation	measure	context type	feature weighting
meronym	cosWeeds	dep	ppmi
	Weeds	dep / joint	ppmi
	ClarkeDE	dep / joint	ppmi / freq
attribute	APSyn	joint	freq
	cosine	joint	freq
	Lin	dep	ppmi
	cosine	dep	ppmi
antonym	SLQS	-	-
	SLQS_row	joint	(freq/ppmi/plmi)
synonym	SLQS_row/SLQS_row_sub	dep	ppmi
	invCL	win2/5/5d	freq
coordination		-	

**Table 3:** Intersection of datasets’ top-performing measures when discriminating between hypernymy and each other relation.

in few other cases, always in combination with syntactic contexts.

Finally, the results show that there is no single combination of measure and parameters that performs consistently well for all datasets and classification tasks. In the following section we analyze the best combination of measure, context type and feature weighting to distinguish hypernymy from any other relation.

## 5.2 Best Measure Per Classification Task

We considered all relations that occurred in two datasets. For such relation, for each dataset, we ranked the measures by their AP@100 score, selecting those with score  $\geq 0.8$ .<sup>6</sup> Table 3 displays the intersection of the datasets’ best measures.

**Hypernym vs. Meronym** The inclusion hypothesis seems to be most effective in discriminating between hypernyms and meronyms under syntactic contexts. We conjecture that the window-based contexts are less effective since they capture topical context words, that might be shared also among holonyms and their meronyms (e.g. *car* will occur with many of the neighbors of *wheel*). However, since meronyms and holonyms often have different functions, their functional contexts, which are expressed in the syntactic context-types, are less shared. This is where they mostly differ from hyponym-hypernym pairs, which are of the same function (e.g. *cat* is a type of *animal*).

Table 2 shows that SLQS performs well in this task on BLESS. This is contrary to previous findings that suggested that SLQS is weak in discriminating between hypernyms and meronyms, as in many cases the holonym is more general than the meronym (Shwartz et al., 2016).<sup>7</sup> The

<sup>6</sup>We considered at least 10 measures, allowing scores slightly lower than 0.8 when others were unavailable.

<sup>7</sup>In the hypernymy dataset of Shwartz et al. (2016),

surprising result could be explained by the nature of meronymy in this dataset: most holonyms in BLESS are rather specific words.

BLESS was built starting from 200 basic level concepts (e.g. *goldfish*) used as the  $x$  words, to which  $y$  words in different relations were associated (e.g. *eye*, for meronymy; *animal*, for hypernymy).  $x$  words represent hyponyms in the hyponym-hypernym pairs, and should therefore not be too general. Indeed, SLQS assigns high scores to hyponym-hypernym pairs. At the same time, in the meronymy relation in BLESS,  $x$  is the holonym and  $y$  is the meronym. For consistency with EVALution, we switched those pairs in BLESS, placing the meronym in the  $x$  slot and the holonym in the  $y$  slot. As a consequence, after the switching, holonyms in BLESS are usually rather specific words (e.g., there are no holonyms like *animal* and *vehicle*, as these words were originally in the  $y$  slot). In most cases, they are not more general than their meronyms ((*eye*, *goldfish*)), yielding low SLQS scores which are easy to separate from hypernyms. We note that this is a weakness of the BLESS dataset, rather than a strength of the measure. For instance, on EVALution, SLQS performs worse (ranked only as high as 13th), as this dataset has no such restriction on the basic level concepts, and may contain pairs like (*eye*, *animal*).

**Hypernym vs. Attribute** Symmetric similarity measures computed on syntactic contexts succeed to discriminate between hypernyms and attributes. Since attributes are syntactically different from hypernyms (in attributes,  $y$  is an adjective), it is unsurprising that they occur in different syntactic contexts, yielding low similarity scores.

nearly 50% of the SLQS false positive pairs were meronym-holonym pairs, in many of which the holonym is more general than the meronym by definition, e.g. (*mauritus*, *africa*).

dataset	hyper vs. relation	best supervised				best unsupervised			
		method	vectors	penalty	AP @100	measure	context type	feature weighting	AP @100
EVALution	meronym	concat	dep-based	$L_2$	0.998	APSyn	joint	freq	0.886
	attribute	concat	Glove-100	$L_2$	1.000	invCL	dep	ppmi	0.877
	antonym	concat	dep-based	$L_2$	1.000	invCL	joint	ppmi	0.773
	synonym	concat	dep-based	$L_1$	0.996	SLQS <sub>sub</sub>	win2	plmi	0.813
BLESS	meronym	concat	Glove-50	$L_1$	1.000	SLQS <sub>sub</sub>	win5	freq	0.939
	coord	concat	Glove-300	$L_1$	1.000	SLQS <sub>row<sub>sub</sub></sub>	joint	plmi	0.938
	attribute	concat	Glove-100	$L_1$	1.000	SLQS <sub>sub</sub>	dep	freq	0.938
	event	concat	Glove-100	$L_1$	1.000	SLQS <sub>sub</sub>	dep	freq	0.847
	random-n	concat	word2vec	$L_1$	0.995	cosWeeds	win2d	ppmi	0.818
	random-j	concat	Glove-200	$L_1$	1.000	SLQS <sub>sub</sub>	dep	freq	0.917
random-v	concat	word2vec	$L_1$	1.000	SLQS <sub>sub</sub>	dep	freq	0.895	
Lenci/ Benotto	antonym	concat	dep-based	$L_2$	0.917	invCL	joint	ppmi	0.807
	synonym	concat	Glove-300	$L_1$	0.946	invCL	win5d	freq	0.914
Weeds	coord	concat	dep-based	$L_2$	0.873	invCL	win2d	freq	0.824
						SLQS <sub>row<sub>sub</sub></sub>	joint	ppmi	

**Table 4:** Best performance on the validation set (10%) of each dataset for the supervised and unsupervised measures, in terms of Average Precision (AP) at  $k = 100$ , for hypernym vs. each single relation.

**Hypernym vs. Antonym** In all our experiments, antonyms were the hardest to distinguish from hypernyms, yielding the lowest performance. We found that SLQS performed reasonably well in this setting. However, the measure variations, context types and feature weightings were not consistent across datasets. SLQS relies on the assumption that  $y$  is a more general word than  $x$ , which is not true for antonyms, making it the most suitable measure for this setting.

**Hypernym vs. Synonym** SLQS performs well also in discriminating between hypernyms and synonyms, in which  $y$  is also not more general than  $x$ . We observed that in the `joint` context type, the difference in SLQS scores between synonyms and hypernyms was the largest. This may stem from the restrictiveness of this context type. For instance, among the most salient contexts we would expect to find informative contexts like *drinks milk* for *cat* and less informative ones like *drinks water* for *animal*, whereas the non-restrictive single dependency context *drinks* would probably be present for both.

Another measure that works well is `invCL`: interestingly, other inclusion-based measures assign high scores to  $(x, y)$  when  $y$  includes many of  $x$ 's contexts, which might be true also for synonyms (e.g. *elevator* and *lift* share many contexts). `invCL`, on the other hand, reduces with the ratio of  $y$ 's contexts included in  $x$ , yielding lower scores for synonyms.

**Hypernym vs. Coordination** We found no consistency among BLESS and Weeds. On Weeds,

inclusion-based measures (`ClarkeDE`, `invCL` and `Weeds`) showed the best results. The best performing measures on BLESS, however, were variants of SLQS, that showed to perform well in cases where the negative relation is symmetric (antonym, synonym and coordination). The difference could be explained by the nature of the datasets: the BLESS test set contains 1,185 hypernymy pairs, with only 129 distinct  $ys$ , many of which are general words like *animal* and *object*. The Weeds test set, on the other hand, was intentionally constructed to contain an overall unique  $y$  in each pair, and therefore contains much more specific  $ys$  (e.g. (*quirk*, *strangeness*)). For this reason, generality-based measures perform well on BLESS, and struggle with Weeds, which is handled better using inclusion-based measures.

### 5.3 Comparison to State-of-the-art Supervised Methods

For comparison with the state-of-the-art, we evaluated several supervised hypernymy detection methods, based on the word embeddings of  $x$  and  $y$ : concatenation  $\vec{v}_x \oplus \vec{v}_y$  (Baroni et al., 2012), difference  $\vec{v}_y - \vec{v}_x$  (Weeds et al., 2014), and ASYM (Roller et al., 2014). We downloaded several pre-trained embeddings (Mikolov et al., 2013; Pennington et al., 2014; Levy and Goldberg, 2014), and trained a logistic regression classifier to predict hypernymy. We used the 90% portion (originally the test set) as the train set, and the other 10% (originally the validation set) as a test set, reporting the best results among different vectors,

	method	AP@100 original	AP@100 switched	$\Delta$
supervised	concat, word2vec, $L_1$	0.995	0.575	-0.42
unsupervised	cosWeeds, win2d, ppmi	0.818	0.882	+0.064

**Table 5:** Average Precision (AP) at  $k = 100$  of the best supervised and unsupervised methods for hypernym vs. random-n, on the original BLESS validation set and the validation set with the artificially added switched hypernym pairs.

method and regularization factor.<sup>8</sup>

Table 4 displays the performance of the best classifier on each dataset, in a hypernym vs. a single relation setting. We also re-evaluated the unsupervised measures, this time reporting the results on the validation set (10%) for comparison.

The overall performance of the embedding-based classifiers is almost perfect, and in particular the best performance is achieved using the concatenation method (Baroni et al., 2012) with either GloVe (Pennington et al., 2014) or the dependency-based embeddings (Levy and Goldberg, 2014). As expected, the unsupervised measures perform worse than the embedding-based classifiers, though generally not bad on their own.

These results may suggest that unsupervised methods should be preferred only when no training data is available, leaving all the other cases to supervised methods. This is, however, not completely true. As others previously noticed, supervised methods do not actually learn the relation between  $x$  and  $y$ , but rather separate properties of either  $x$  or  $y$ . Levy et al. (2015b) named this the “lexical memorization” effect, i.e. memorizing that certain  $y$ s tend to appear in many positive pairs (*prototypical hypernyms*).

On that account, the Weeds dataset has been designed to avoid such memorization, with every word occurring once in each slot of the relation. While the performance of the supervised methods on this dataset is substantially lower than their performance on other datasets, it is yet well above the random baseline which we might expect from a method that can only memorize words it has seen during training.<sup>9</sup> This is an indication that supervised methods can abstract away from the words.

Indeed, when we repeated the experiment with a lexical split of each dataset, i.e., such that the train and test set consist of distinct vocabularies, we found that the supervised methods’ performance did not decrease dramatically, in contrast to the

findings of Levy et al. (2015b). The large performance gaps reported by Levy et al. (2015b) might be attributed to the size of their training sets. Their lexical split discarded around half of the pairs in the dataset and split the rest of the pairs equally to train and test, resulting in a relatively small train set. We performed the split such that only around 30% of the pairs in each dataset were discarded, and split the train and test sets with a ratio of roughly 90/10%, obtaining large enough train sets.

Our experiment suggests that rather than memorizing the verbatim *prototypical hypernyms*, the supervised models might learn that certain regions in the vector space pertain to prototypical hypernyms. For example, *device* (from the BLESS train set) and *appliance* (from the BLESS test set) are two similar words, which are both prototypical hypernyms. Another interesting observation was recently made by Roller and Erk (2016): they showed that when dependency-based embeddings are used, supervised distributional methods trace  $x$  and  $y$ ’s separate occurrences in different slots of Hearst patterns (Hearst, 1992).

Whether supervised methods only memorize or also learn, it is more consensual that they lack the ability to capture the relation between  $x$  and  $y$ , and that they rather indicate how likely  $y$  ( $x$ ) is to be a hypernym (hyponym) (Levy et al., 2015b; Santus et al., 2016a; Shwartz et al., 2016; Roller and Erk, 2016). While this information is valuable, it cannot be solely relied upon for classification.

To better understand the extent of this limitation, we conducted an experiment in a similar manner to the switched hypernym pairs in Santus et al. (2016a). We used BLESS, which is the only dataset with random pairs. For each hypernym pair  $(x_1, y_1)$ , we sampled a word  $y_2$  that participates in another hypernym pair  $(x_2, y_2)$ , such that  $(x_1, y_2)$  is not in the dataset, and added  $(x_1, y_2)$  as a random pair. We added 139 new pairs to the validation set, such as (*rifle, animal*) and (*salmon, weapon*). We then used the best supervised and unsupervised methods for hypernym vs. random-n on BLESS to re-classify the revised validation set. Table 5 displays the experiment results.

<sup>8</sup>In our preliminary experiments we also trained other classifiers used in the distributional hypernymy detection literature (SVM and SVM+RBF kernel), that performed similarly. We report the results for logistic regression, since we use the prediction probabilities to measure average precision.

<sup>9</sup>The dataset is balanced between its two classes.



The switched hypernym experiment paints a much less optimistic picture of the embeddings’ actual performance, with a drop of 42 points in average precision. 121 out of the 139 switched hypernym pairs were falsely classified as hypernyms. Examining the  $y$  words of these pairs reveals general words that appear in many hypernym pairs (e.g. *animal*, *object*, *vehicle*). The unsupervised measure was not similarly affected by the switched pairs, and the performance even slightly increased. This result is not surprising, since most unsupervised measures aim to capture aspects of the relation between  $x$  and  $y$ , while not relying on information about one of the words in the pair.<sup>10</sup>

## 6 Discussion

The results in Section 5 suggest that a supervised method using the unsupervised measures as features could possibly be the best of both worlds. We would expect it to be more robust than embedding-based methods on the one hand, while being more informative than any single unsupervised measure on the other hand.

Such a method was developed by Santus et al. (2016a), however using mostly features that describe a single word, e.g. frequency and entropy. It was shown to be competitive with the state-of-the-art supervised methods. With that said, it was also shown to be sensitive to the distribution of training examples in a specific dataset, like the embedding-based methods.

We conducted a similar experiment, with a much larger number of unsupervised features, namely the various measure scores, and encountered the same issue. While the performance was good, it dropped dramatically when the model was tested on a different test set.

We conjecture that the problem stems from the currently available datasets, which are all somewhat artificial and biased. Supervised methods which are strongly based on the relation between the words, e.g. those that rely on path-based information (Shwartz et al., 2016), manage to overcome the bias. Distributional methods, on the other hand, are based on a weaker notion of the relation between words, hence are more prone to overfit the distribution of training instances in a specific dataset. In the future, we hope that new

<sup>10</sup>Turney and Mohammad (2015) have also shown that unsupervised methods are more robust than supervised ones in a transfer-learning experiment, when the “training data” was used to tune their parameters.

datasets will be available for the task, which would be drawn from corpora and will reflect more realistic distributions of words and semantic relations.

## 7 Conclusion

We performed an extensive evaluation of unsupervised methods for discriminating hypernyms from other semantic relations. We found that there is no single combination of measure and parameters which is always preferred; however, we suggested a principled linguistic-based analysis of the most suitable measure for each task that yields consistent performance across different datasets.

We investigated several new variants of existing methods, and found that some variants of SLQS turned out to be superior on certain tasks. In addition, we have tested for the first time the `joint` context type (Chersoni et al., 2016), which was found to be very discriminative, and might hopefully benefit other semantic tasks.

For comparison, we evaluated the state-of-the-art supervised methods on the datasets, and they have shown to outperform the unsupervised ones, while also being efficient and easier to use. However, a deeper analysis of their performance demonstrated that, as previously suggested, these methods do not capture the relation between  $x$  and  $y$ , but rather indicate the “prior probability” of either word to be a hyponym or a hypernym. As a consequence, supervised methods are sensitive to the distribution of examples in a particular dataset, making them less reliable for real-world applications. Being motivated by linguistic hypotheses, and independent from training data, unsupervised measures were shown to be more robust. In this sense, unsupervised methods can still play a relevant role, especially if combined with supervised methods, in the decision whether the relation holds or not.

## Acknowledgments

The authors would like to thank Ido Dagan, Alessandro Lenci, and Yuji Matsumoto for their help and advice. Vered Shwartz is 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). Enrico Santus is partially supported by HK PhD Fellowship Scheme under PF12-13656.

## References

- Marco Baroni and Alessandro Lenci. 2010. Distributional memory: A general framework for corpus-based semantics. *Computational Linguistics*, 36(4):673–721.
- Marco Baroni and Alessandro Lenci. 2011. Proceedings of the gems 2011 workshop on geometrical models of natural language semantics. In *How we BLESSed distributional semantic evaluation*, pages 1–10. Association for Computational Linguistics.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. The wacky wide web: a collection of very large linguistically processed web-crawled corpora. *Language resources and evaluation*, 43(3):209–226.
- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 23–32. Association for Computational Linguistics.
- Giulia Benotto. 2015. Distributional models for semantic relations: A study on hyponymy and antonymy. *PhD Thesis, University of Pisa*.
- John A Bullinaria and Joseph P Levy. 2007. Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior research methods*, 39(3):510–526.
- Emmanuele Chersoni, Enrico Santus, Alessandro Lenci, Philippe Blache, and Chu-Ren Huang. 2016. Representing verbs with rich contexts: an evaluation on verb similarity. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1967–1972. Association for Computational Linguistics.
- Kenneth Ward Church and Patrick Hanks. 1990. Word association norms, mutual information, and lexicography. *Computational linguistics*, 16(1):22–29.
- Daoud Clarke. 2009. Context-theoretic semantics for natural language: an overview. In *Proceedings of the Workshop on Geometrical Models of Natural Language Semantics*, pages 112–119. Association for Computational Linguistics.
- Ido Dagan, Dan Roth, and Mark Sammons. 2013. *Recognizing textual entailment*. Morgan & Claypool Publishers.
- Stefan Evert. 2005. The statistics of word cooccurrences: word pairs and collocations. *Dissertation*.
- Stefan Evert. 2008. Corpora and collocations. *Corpus linguistics. An international handbook*, 2:223–233.
- Adriano Ferraresi. 2007. Building a very large corpus of english obtained by web crawling: ukwac. *Masters thesis, University of Bologna, Italy*.
- Zellig S. Harris. 1954. Distributional structure. *Word*, 10(2-3):146–162.
- Marti A. Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *COLING 1992 Volume 2: The 15th International Conference on Computational Linguistics*.
- Douwe Kiela, Laura Rimell, Ivan Vulić, and Stephen Clark. 2015. Exploiting image generality for lexical entailment detection. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 119–124. Association for Computational Linguistics.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2010. Directional distributional similarity for lexical inference. *Natural Language Engineering*, 16(04):359–389.
- Alessandro Lenci and Giulia Benotto. 2012. Identifying hypernyms in distributional semantic spaces. In *\*SEM 2012: 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 (SemEval 2012)*, pages 75–79. Association for Computational Linguistics.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 302–308. Association for Computational Linguistics.
- Omer Levy, Yoav Goldberg, and Ido Dagan. 2015a. Improving distributional similarity with lessons learned from word embeddings. *Transactions of the Association for Computational Linguistics*, 3:211–225.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015b. Do supervised distributional methods really learn lexical inference relations? In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 970–976. Association for Computational Linguistics.
- Dekang Lin. 1998. An information-theoretic definition of similarity. *ICML*, 98:296–304.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositional-ity. In *Advances in neural information processing systems*, pages 3111–3119.
- Roberto Navigli, Paola Velardi, and Stefano Faralli. 2011. A graph-based algorithm for inducing lexical taxonomies from scratch. In *Proceedings of*

- the Twenty-Second International Joint Conference on Artificial Intelligence*, volume 11, pages 1872–1877.
- Sebastian Padó and Mirella Lapata. 2007. Dependency-based construction of semantic space models. *Computational Linguistics*, 33(2):161–199.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543. Association for Computational Linguistics.
- Laura Rimell. 2014. Distributional lexical entailment by topic coherence. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 511–519. Association for Computational Linguistics.
- Stephen Roller and Katrin Erk. 2016. Relations such as hypernymy: Identifying and exploiting hearst patterns in distributional vectors for lexical entailment. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2163–2172. Association for Computational Linguistics.
- Stephen Roller, Katrin Erk, and Gemma Boleda. 2014. Inclusive yet selective: Supervised distributional hypernymy detection. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 1025–1036. Dublin City University and Association for Computational Linguistics.
- Gerard Salton and Michael J. McGill. 1986. *Introduction to modern information retrieval*. McGraw-Hill, Inc.
- 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. Association for Computational Linguistics.
- Enrico Santus, Frances Yung, Alessandro Lenci, and Chu-Ren Huang. 2015. Proceedings of the 4th workshop on linked data in linguistics: Resources and applications. In *EVALution 1.0: an Evolving Semantic Dataset for Training and Evaluation of Distributional Semantic Models*, pages 64–69. Association for Computational Linguistics.
- Enrico Santus, Alessandro Lenci, Tin-Shing Chiu, Qin Lu, and Chu-Ren Huang. 2016a. Nine features in a random forest to learn taxonomical semantic relations. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*, pages 4557–4564. European Language Resources Association (ELRA).
- Enrico Santus, Alessandro Lenci, Tin-Shing Chiu, Qin Lu, and Chu-Ren Huang. 2016b. Unsupervised measure of word similarity: How to outperform co-occurrence and vector cosine in vsms. In *Thirtieth AAAI Conference on Artificial Intelligence*, pages 4260–4261.
- Vered Shwartz, Yoav Goldberg, and Ido Dagan. 2016. Improving hypernymy detection with an integrated path-based and distributional method. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2389–2398. Association for Computational Linguistics.
- Rion Snow, Daniel Jurafsky, and Andrew Y. Ng. 2005. Learning syntactic patterns for automatic hypernym discovery. In *Advances in Neural Information Processing Systems 17*, pages 1297–1304. MIT Press.
- 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 44th Annual Meeting of the Association for Computational Linguistics*, pages 801–808. Association for Computational Linguistics.
- Peter D. Turney and Saif M. Mohammad. 2015. Experiments with three approaches to recognizing lexical entailment. *Natural Language Engineering*, 21(03):437–476.
- Julie Weeds and David Weir. 2003. A general framework for distributional similarity. In *Proceedings of the 2003 Conference on Empirical Methods in Natural Language Processing*, pages 81–88.
- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. Learning to distinguish hypernyms and co-hyponyms. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2249–2259. Dublin City University and Association for Computational Linguistics.