

Hearst Patterns Revisited: Automatic Hypernym Detection from Large Text Corpora

Stephen Roller, Douwe Kiela, and Maximilian Nickel

Facebook AI Research

{roller, dkiela, maxn}@fb.com

Abstract

Methods for unsupervised hypernym detection may broadly be categorized according to two paradigms: pattern-based and distributional methods. In this paper, we study the performance of both approaches on several hypernymy tasks and find that simple pattern-based methods consistently outperform distributional methods on common benchmark datasets. Our results show that pattern-based models provide important contextual constraints which are not yet captured in distributional methods.

1 Introduction

Hierarchical relationships play a central role in knowledge representation and reasoning. Hypernym detection, i.e., the modeling of word-level hierarchies, has long been an important task in natural language processing. Starting with Hearst (1992), pattern-based methods have been one of the most influential approaches to this problem. Their key idea is to exploit certain *lexico-syntactic patterns* to detect *is-a* relations in text. For instance, patterns like “NP_y such as NP_x”, or “NP_x and other NP_y” often indicate hypernymy relations of the form *x is-a y*. Such patterns may be predefined, or they may be learned automatically (Snow et al., 2004; Shwartz et al., 2016). However, a well-known problem of Hearst-like patterns is their extreme sparsity: words must co-occur in exactly the right configuration, or else no relation can be detected.

To alleviate the sparsity issue, the focus in hypernymy detection has recently shifted to distributional representations, wherein words are represented as vectors based on their distribution across large corpora. Such methods offer rich representations of lexical meaning, alleviating the sparsity problem, but require specialized similarity mea-

asures to distinguish different lexical relationships. The most successful measures to date are generally inspired by the Distributional Inclusion Hypothesis (DIH) (Zhitomirsky-Geffet and Dagan, 2005), which states roughly that contexts in which a narrow term *x* may appear (“cat”) should be a subset of the contexts in which a broader term *y* (“animal”) may appear. Intuitively, the DIH states that we should be able to replace any occurrence of “cat” with “animal” and still have a valid utterance. An important insight from work on distributional methods is that the definition of context is often critical to the success of a system (Shwartz et al., 2017). Some distributional representations, like positional or dependency-based contexts, may even capture crude Hearst pattern-like features (Levy et al., 2015; Roller and Erk, 2016).

While both approaches for hypernym detection rely on co-occurrences within certain contexts, they differ in their context selection strategy: pattern-based methods use predefined manually-curated patterns to generate high-precision extractions while DIH methods rely on unconstrained word co-occurrences in large corpora.

Here, we revisit the idea of using pattern-based methods for hypernym detection. We evaluate several pattern-based models on modern, large corpora and compare them to methods based on the DIH. We find that simple pattern-based methods consistently outperform specialized DIH methods on several difficult hypernymy tasks, including detection, direction prediction, and graded entailment ranking. Moreover, we find that taking low-rank embeddings of pattern-based models substantially improves performance by remedying the sparsity issue. Overall, our results show that Hearst patterns provide high-quality and robust predictions on large corpora by capturing important contextual constraints, which are not yet modeled in distributional methods.

2 Models

In the following, we discuss pattern-based and distributional methods to detect hypernymy relations. We explicitly consider only relatively simple pattern-based approaches that allow us to directly compare their performance to DIH-based methods.

2.1 Pattern-based Hypernym Detection

First, let $\mathcal{P} = \{(x, y)\}_{i=1}^n$ denote the set of hypernymy relations that have been extracted via Hearst patterns from a text corpus \mathcal{T} . Furthermore let $w(x, y)$ denote the count of how often (x, y) has been extracted and let $W = \sum_{(x,y) \in \mathcal{P}} w(x, y)$ denote the total number extractions. In the first, most direct application of Hearst patterns, we then simply use the counts $w(x, y)$ or, equivalently, the extraction probability

$$p(x, y) = \frac{w(x, y)}{W} \quad (1)$$

to predict hypernymy relations from \mathcal{T} .

However, simple extraction probabilities as in Equation (1) are skewed by the occurrence probabilities of their constituent words. For instance, it is more likely that we extract *(France, country)* over *(France, republic)*, just because the word *country* is more likely to occur than *republic*. This skew in word distributions is well-known for natural language and also translates to Hearst patterns (see also Figure 1). For this reason, we also consider predicting hypernymy relations based on the Pointwise Mutual Information of Hearst patterns: First, let $p^-(x) = \sum_{(x,y) \in \mathcal{P}} w(x, y)/W$ and $p^+(x) = \sum_{(y,x) \in \mathcal{P}} w(y, x)/W$ denote the probability that x occurs as a hyponym and hypernym, respectively. We then define the Positive Pointwise Mutual Information for (x, y) as

$$\text{ppmi}(x, y) = \max \left(0, \log \frac{p(x, y)}{p^-(x)p^+(y)} \right). \quad (2)$$

While Equation (2) can correct for different word occurrence probabilities, it cannot handle missing data. However, sparsity is one of the main issues when using Hearst patterns, as a necessarily incomplete set of extraction rules will lead inevitably to missing extractions. For this purpose, we also study low-rank embeddings of the PPMI matrix, which allow us to make predictions for unseen pairs. In particular, let $m = |\{x : (x, y) \in \mathcal{P} \vee (y, x) \in \mathcal{P}\}|$ denote the number of unique terms in \mathcal{P} . Furthermore, let $X \in \mathbb{R}^{m \times m}$ be the PPMI matrix with




Figure 1: Frequency distribution of words appearing in Hearst patterns.

entries $M_{xy} = \text{ppmi}(x, y)$ and let $M = U\Sigma V^\top$ be its *Singular Value Decomposition* (SVD). We can then predict hypernymy relations based on the truncated SVD of M via

$$\text{spmi}(x, y) = \mathbf{u}_x^\top \Sigma_r \mathbf{v}_y \quad (3)$$

where $\mathbf{u}_x, \mathbf{v}_y$ denote the x -th and y -th row of U and V , respectively, and where Σ_r is the diagonal matrix of truncated singular values (in which all but the r largest singular values are set to zero).

Equation (3) can be interpreted as a smoothed version of the observed PPMI matrix. Due to the truncation of singular values, Equation (3) computes a low-rank embedding of M where similar words (in terms of their Hearst patterns) have similar representations. Since Equation (3) is defined for all pairs (x, y) , it allows us to make hypernymy predictions based on the similarity of words. We also consider factorizing a matrix that is constructed from occurrence probabilities as in Equation (1), denoted by $\text{sp}(x, y)$. This approach is then closely related to the method of [Cederberg and Widdows \(2003\)](#), which has been proposed to improve precision and recall for hypernymy detection from Hearst patterns.

2.2 Distributional Hypernym Detection

Most unsupervised distributional approaches for hypernymy detection are based on variants of the Distributional Inclusion Hypothesis ([Weeds et al., 2004](#); [Kotlerman et al., 2010](#); [Santus et al., 2014](#); [Lenci and Benotto, 2012](#); [Shwartz et al., 2017](#)). Here, we compare to two methods with strong empirical results. As with most DIH measures, they are only defined for large, sparse, positively-valued distributional spaces. First, we consider *WeedsPrec* ([Weeds et al., 2004](#)) which captures the features of

\mathbf{x} which are *included* in the set of a broader term’s features, \mathbf{y} :

$$\text{WeedsPrec}(\mathbf{x}, \mathbf{y}) = \frac{\sum_{i=1}^n x_i * \mathbb{1}_{y_i > 0}}{\sum_{i=1}^n x_i}$$

Second, we consider *invCL* (Lenci and Benotto, 2012) which introduces a notion of distributional *exclusion* by also measuring the degree to which the broader term contains contexts *not* used by the narrower term. In particular, let

$$\text{CL}(\mathbf{x}, \mathbf{y}) = \frac{\sum_{i=1}^n \min(x_i, y_i)}{\sum_{i=1}^n x_i}$$

denote the degree of inclusion of x in y as proposed by Clarke (2009). To measure both the inclusion of x in y and the non-inclusion of y in x , *invCL* is then defined as

$$\text{invCL}(\mathbf{x}, \mathbf{y}) = \sqrt{\text{CL}(\mathbf{x}, \mathbf{y}) * (1 - \text{CL}(\mathbf{y}, \mathbf{x}))}$$

Although most unsupervised distributional approaches are based on the DIH, we also consider the distributional SLQS model based on an alternative *informativeness* hypothesis (Santus et al., 2014; Shwartz et al., 2017). Intuitively, the SLQS model presupposes that general words appear mostly in uninformative contexts, as measured by entropy. Specifically, SLQS depends on the median entropy of a term’s top N contexts, defined as

$$E_x = \text{median}_{i=1}^N [H(c_i)],$$

where $H(c_i)$ is the Shannon entropy of context c_i across all terms, and N is chosen in hyperparameter selection. Finally, SLQS is defined using the ratio between the two terms:

$$\text{SLQS}(x, y) = 1 - \frac{E_x}{E_y}.$$

Since the SLQS model only compares the relative generality of two terms, but does not make judgment about the terms’ relatedness, we report SLQS-cos, which multiplies the SLQS measure by cosine similarity of x and y (Santus et al., 2014).

For completeness, we also include cosine similarity as a baseline in our evaluation.

3 Evaluation

To evaluate the relative performance of pattern-based and distributional models, we apply them to several challenging hypernymy tasks.

Pattern

X which is a (example|class|kind|. . .) of Y
 X (and|or) (any|some) other Y
 X which is called Y
 X is JJS (most)? Y
 X a special case of Y
 X is an Y that
 X is a !(member|part|given) Y
 !(features|properties) Y such as X_1, X_2, \dots
 (Unlike|like) (most|all|any|other) Y, X
 Y including X_1, X_2, \dots

Table 1: Hearst patterns used in this study. Patterns are lemmatized, but listed as inflected for clarity.

3.1 Tasks

Detection: In hypernymy detection, the task is to classify whether pairs of words are in a hypernymy relation. For this task, we evaluate all models on five benchmark datasets: First, we employ the noun-noun subset of BLESS, which contains hypernymy annotations for 200 concrete, mostly unambiguous nouns. Negative pairs contain a mixture of co-hyponymy, meronymy, and random pairs. This version contains 14,542 total pairs with 1,337 positive examples. Second, we evaluate on LEDS (Baroni et al., 2012), which consists of 2,770 noun pairs balanced between positive hypernymy examples, and randomly shuffled negative pairs. We also consider EVAL (Santus et al., 2015), containing 7,378 pairs in a mixture of hypernymy, synonymy, antonymy, meronymy, and adjectival relations. EVAL is notable for its absence of random pairs. The largest dataset is SHWARTZ (Shwartz et al., 2016), which was collected from a mixture of WordNet, DBPedia, and other resources. We limit ourselves to a 52,578 pair subset excluding multiword expressions. Finally, we evaluate on WBLESS (Weeds et al., 2014), a 1,668 pair subset of BLESS, with negative pairs being selected from co-hyponymy, random, and hyponymy relations. Previous work has used different metrics for evaluating on BLESS (Lenci and Benotto, 2012; Levy et al., 2015; Roller and Erk, 2016). We chose to evaluate the global ranking using Average Precision. This allowed us to use the same metric on all detection benchmarks, and is consistent with evaluations in Shwartz et al. (2017).

Direction: In direction prediction, the task is to identify which term is broader in a given pair

of words. For this task, we evaluate all models on three datasets described by [Kiela et al. \(2015\)](#): On BLESS, the task is to predict the direction for all 1337 positive pairs in the dataset. Pairs are only counted correct if the hypernymy direction scores higher than the reverse direction, i.e. $\text{score}(x, y) > \text{score}(y, x)$. We reserve 10% of the data for validation, and test on the remaining 90%. On WBLESS, we follow prior work ([Nguyen et al., 2017](#); [Vulić and Mrkšić, 2017](#)) and perform 1000 random iterations in which 2% of the data is used as a validation set to learn a classification threshold, and test on the remainder of the data. We report average accuracy across all iterations. Finally, we evaluate on BIBLESS ([Kiela et al., 2015](#)), a variant of WBLESS with hypernymy and hyponymy pairs explicitly annotated for their direction. Since this task requires three-way classification (hypernymy, hyponymy, and other), we perform two-stage classification. First, a threshold is tuned using 2% of the data, identifying whether a pair exhibits hypernymy in either direction. Second, the relative comparison of scores determines which direction is predicted. As with WBLESS, we report the average accuracy over 1000 iterations.

Graded Entailment: In graded entailment, the task is to quantify the *degree* to which a hypernymy relation holds. For this task, we follow prior work ([Nickel and Kiela, 2017](#); [Vulić and Mrkšić, 2017](#)) and use the noun part of HYPERLEX ([Vulić et al., 2017](#)), consisting of 2,163 noun pairs which are annotated to what degree *x is-a y* holds on a scale of $[0, 6]$. For all models, we report Spearman’s rank correlation ρ . We handle out-of-vocabulary (OOV) words by assigning the median of the scores (computed across the training set) to pairs with OOV words.

3.2 Experimental Setup

Pattern-based models: We extract Hearst patterns from the concatenation of Gigaword and Wikipedia, and prepare our corpus by tokenizing, lemmatizing, and POS tagging using CoreNLP 3.8.0. The full set of Hearst patterns is provided in Table 1. Our selected patterns match prototypical Hearst patterns, like “animals such as cats,” but also include broader patterns like “New Year is the most important holiday.” Leading and following noun phrases are allowed to match limited modifiers (compound nouns, adjectives, etc.), in which case we also generate a hit for the head of the noun phrase. Dur-

ing postprocessing, we remove pairs which were not extracted by at least two distinct patterns. We also remove any pair (y, x) if $p(y, x) < p(x, y)$. The final corpus contains roughly 4.5M matched pairs, 431K unique pairs, and 243K unique terms. For SVD-based models, we select the rank from $r \in \{5, 10, 15, 20, 25, 50, 100, 150, 200, 250, 300, 500, 1000\}$ on the validation set. The other pattern-based models do not have any hyperparameters.

Distributional models: For the distributional baselines, we employ the large, sparse distributional space of [Shwartz et al. \(2017\)](#), which is computed from UkWaC and Wikipedia, and is known to have strong performance on several of the detection tasks. The corpus was POS tagged and dependency parsed. Distributional contexts were constructed from adjacent words in dependency parses ([Padó and Lapata, 2007](#); [Levy and Goldberg, 2014](#)). Targets and contexts which appeared fewer than 100 times in the corpus were filtered, and the resulting co-occurrence matrix was PPMI transformed.¹ The resulting space contains representations for 218K words over 732K context dimensions. For the SLQS model, we selected the number of contexts N from the same set of options as the SVD rank in pattern-based models.

3.3 Results

Table 2 shows the results from all three experimental settings. In nearly all cases, we find that pattern-based approaches substantially outperform all three distributional models. Particularly strong improvements can be observed on BLESS (0.76 average precision vs 0.19) and WBLESS (0.96 vs. 0.69) for the detection tasks and on all directionality tasks. For directionality prediction on BLESS, the SVD models surpass even the state-of-the-art *supervised* model of [Vulić and Mrkšić \(2017\)](#). Moreover, both SVD models perform generally better than their sparse counterparts on all tasks and datasets except on HYPERLEX. We performed a posthoc analysis of the validation sets comparing the ppmi and spmi models, and found that the truncated SVD improved recall via its matrix completion properties. We also found that the spmi model downweighted

¹In addition, we also experimented with further distributional spaces and weighting schemes from [Shwartz et al. \(2017\)](#). We also experimented with distributional spaces using the same corpora and preprocessing as the Hearst patterns (i.e., Wikipedia and Gigaword). We found that the reported setting generally performed best, and omit others for brevity.

	Detection (AP)					Direction (Acc.)			Graded (ρ_s)
	BLESS	EVAL	LEDS	SHWARTZ	WBLESS	BLESS	WBLESS	BIBLESS	HYPERLEX
Cosine	.12	.29	.71	.31	.53	.00	.54	.52	.14
WeedsPrec	.19	.39	.87	.43	.68	.63	.59	.45	.43
invCL	.18	.37	.89	.38	.66	.64	.60	.47	.43
SLQS	.15	.35	.60	.38	.69	.75	.67	.51	.16
p(x, y)	.49	.38	.71	.29	.74	.46	.69	.62	.62
ppmi(x, y)	.45	.36	.70	.28	.72	.46	.68	.61	.60
sp(x, y)	.66	.45	.81	.41	.91	.96	.84	.80	.51
spmi(x, y)	.76	.48	.84	.44	.96	.96	.87	.85	.53

Table 2: Experimental results comparing distributional and pattern-based methods in all settings.

many high-scoring outlier pairs composed of rare terms.

When comparing the $p(x, y)$ and ppmi models to distributional models, we observe mixed results. The SHWARTZ dataset is difficult for sparse models due to its very long tail of low frequency words that are hard to cover using Hearst patterns. On EVAL, Hearst-pattern based methods get penalized by OOV words, due to the large number of verbs and adjectives in the dataset, which are not captured by our patterns. However, in 7 of the 9 datasets, at least one of the sparse models outperforms all distributional measures, showing that Hearst patterns can provide strong performance on large corpora.

4 Conclusion

We studied the relative performance of Hearst pattern-based methods and DIH-based methods for hypernym detection. Our results show that the pattern-based methods substantially outperform DIH-based methods on several challenging benchmarks. We find that embedding methods alleviate sparsity concerns of pattern-based approaches and substantially improve coverage. We conclude that Hearst patterns provide important contexts for the detection of hypernymy relations that are not yet captured in DIH models. Our code is available at <https://github.com/facebookresearch/hypernysuite>.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful suggestions. We also thank Vered Shwartz, Enrico Santus, and Dominik Schlechtweg for providing us with their distributional spaces and baseline implementations.

References

- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *Proceedings of the 2012 Conference of the European Chapter of the Association for Computational Linguistics*, pages 23–32, Avignon, France.
- Scott Cederberg and Dominic Widdows. 2003. Using Isa and noun coordination information to improve the recall and precision of automatic hyponymy extraction. In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*.
- Daoud Clarke. 2009. Context-theoretic semantics for natural language: An overview. In *Proceedings of the 2011 Workshop on GEometrical Models of Natural Language Semantics*, pages 112–119, Athens, Greece. Association for Computational Linguistics.
- Marti A. Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proceedings of the 1992 Conference on Computational Linguistics*, pages 539–545, Nantes, France.
- 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, Beijing, China. Association for Computational Linguistics.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2010. Directional distributional similarity for lexical inference. *Natural Language Engineering*, 16:359–389.
- Alessandro Lenci and Giulia Benotto. 2012. Identifying hypernoms in distributional semantic spaces. In *Proceedings of the First Joint Conference on Lexical and Computational Semantics*, pages 75–79, Montréal, Canada. 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*, pages 302–308, Baltimore, Maryland. Association for Computational Linguistics.
- Omer Levy, Steffen Remus, Chris Biemann, and Ido Dagan. 2015. Do supervised distributional methods really learn lexical inference relations? In *Proceedings of the 2015 North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 970–976, Denver, Colorado.
- Kim Anh Nguyen, Maximilian Keper, Sabine Schulte im Walde, and Ngoc Thang Vu. 2017. Hierarchical Embeddings for Hypernymy Detection and Directionality. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 233–243, Copenhagen, Denmark.
- Maximillian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6338–6347. Curran Associates, Inc.
- Sebastian Padó and Mirella Lapata. 2007. Dependency-based construction of semantic space models. *Computational Linguistics*, 33(2):161–199.
- 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*, Austin, Texas, USA. Association for Computational Linguistics.
- Enrico Santus, Alessandro Lenci, Qin Lu, and Sabine Schulte im Walde. 2014. Chasing hypernoms in vector spaces with entropy. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics, volume 2: Short Papers*, pages 38–42, Gothenburg, Sweden. Association for Computational Linguistics.
- Enrico Santus, Frances Yung, Alessandro Lenci, and Chu-Ren Huang. 2015. EVALution 1.0: An evolving semantic dataset for training and evaluation of distributional semantic models. In *Proceedings of the Fourth Workshop on Linked Data in Linguistics: Resources and Applications*, pages 64–69, Beijing, China. Association for Computational Linguistics.
- 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*, pages 2389–2398, Berlin, Germany. Association for Computational Linguistics.
- Vered Shwartz, Enrico Santus, and Dominik Schlechtweg. 2017. Hypernoms under siege: Linguistically-motivated artillery for hypernymy detection. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 65–75, Valencia, Spain. Association for Computational Linguistics.
- Rion Snow, Daniel Jurafsky, and Andrew Y. Ng. 2004. Learning syntactic patterns for automatic hypernym discovery. In L. K. Saul, Y. Weiss, and L. Bottou, editors, *Advances in Neural Information Processing Systems 17*, pages 1297–1304. MIT Press.
- Ivan Vulić, Daniela Gerz, Douwe Kiela, Felix Hill, and Anna Korhonen. 2017. Hyperlex: A large-scale evaluation of graded lexical entailment. *Computational Linguistics*, 43(4):781–835.
- Ivan Vulić and Nikola Mrkšić. 2017. Specialising word vectors for lexical entailment. *CoRR*, abs/1710.06371.
- Julie Weeds, Daoud Clarke, Jeremy Reffin, David Weir, and Bill Keller. 2014. Learning to distinguish hypernoms and co-hyponyms. In *Proceedings of the 2014 International Conference on Computational Linguistics*, pages 2249–2259, Dublin, Ireland.
- Julie Weeds, David Weir, and Diana McCarthy. 2004. Characterising measures of lexical distributional similarity. In *Proceedings of the 2004 International Conference on Computational Linguistics*, pages 1015–1021, Geneva, Switzerland.
- Maayan Zhitomirsky-Geffet and Ido Dagan. 2005. The distributional inclusion hypotheses and lexical entailment. In *Proceedings of the 2005 Annual Meeting of the Association for Computational Linguistics*, pages 107–114, Ann Arbor, Michigan.