

Robust Cross-lingual Hypernymy Detection using Dependency Context

Shyam Upadhyay^{1*}, Yogarshi Vyas^{2*}, Marine Carpuat², Dan Roth¹

¹ Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA

² Department of Computer Science, University of Maryland, College Park, MD

shyamupa@seas.upenn.edu, yogarshi@cs.umd.edu,
marine@cs.umd.edu, danroth@seas.upenn.edu

Abstract

Cross-lingual Hypernymy Detection involves determining if a word in one language (“fruit”) is a hypernym of a word in another language (“pomme” i.e. apple in French). The ability to detect hypernymy cross-lingually can aid in solving cross-lingual versions of tasks such as textual entailment and event coreference. We propose BiSPARSE-DEP, a family of unsupervised approaches for cross-lingual hypernymy detection, which learns sparse, bilingual word embeddings based on dependency contexts. We show that BiSPARSE-DEP can significantly improve performance on this task, compared to approaches based only on lexical context. Our approach is also robust, showing promise for low-resource settings: our dependency-based embeddings can be learned using a parser trained on related languages, with negligible loss in performance. We also crowd-source a challenging dataset for this task on four languages – Russian, French, Arabic, and Chinese. Our embeddings and datasets are publicly available.¹

1 Introduction

Translation helps identify correspondences in bilingual texts, but other asymmetric semantic relationships can improve language understanding when translations are not exactly equivalent. One such relationship is *cross-lingual hypernymy* – identifying that *écureuil* (“squirrel” in French) is a kind of *rodent*, or *ворона* (“crow” in Russian) is a kind of *bird*. The ability to detect hypernyms across languages serves as a building block in a range of cross-lingual tasks, including Recognizing Textual Entailment (RTE) (Negri et al., 2012,

2013), constructing multilingual taxonomies (Fu et al., 2014), event coreference across multilingual news sources (Vossen et al., 2015), and evaluating Machine Translation output (Padó et al., 2009).

Building models that can robustly identify hypernymy across the spectrum of human languages is a challenging problem, that is further compounded in low resource settings. At first glance, translating words to English and then identifying hypernyms in a monolingual setting may appear to be a sufficient solution. However, this approach cannot capture many phenomena. For instance, the English words *cook*, *leader* and *supervisor* can all be hypernyms of the French word *chef*, as the French word does not have an exact translation in English covering its possible usages. However, translating *chef* to *cook* and then determining hypernymy monolingually precludes identifying *leader* or *supervisor* as hypernyms of *chef*. Similarly, language-specific usage patterns can also influence hypernymy decisions. For instance, the French word *chroniqueur* translates to *chronicler* in English, but is more frequently used in French to refer to journalists (making *journalist* its hypernym).²

This motivates approaches that *directly* detect hypernymy in the cross-lingual setting by extending distributional methods for detecting monolingual hypernymy, as in our prior work (Vyas and Carpuat, 2016). State-of-the-art distributional approaches (Roller and Erk, 2016; Shwartz et al., 2017) for detecting monolingual hypernymy require syntactic analysis (eg. dependency parsing), which may not be available for many languages. Additionally, limited training resources make unsupervised methods more desirable than supervised hypernymy detection approaches (Roller and Erk,

* These authors contributed equally.

¹<https://github.com/yogarshi/bisparsedep/>

²All examples are from our dataset.

2016). Furthermore, monolingual distributional approaches cannot be applied directly to the cross-lingual task, because the vector spaces of two languages need to be aligned using a cross-lingual resource (a bilingual dictionary, for instance).

We tackle these challenges by proposing BISPARSE-DEP - a family of robust, unsupervised approaches for identifying cross-lingual hypernymy. BISPARSE-DEP uses a cross-lingual word embedding model learned from a small bilingual dictionary and a variety of monolingual syntactic context extracted from a dependency parsed corpus. BISPARSE-DEP exhibits robust behavior along multiple dimensions. In the absence of a dependency treebank for a language, it can learn embeddings using a parser trained on related languages. When exposed to less monolingual data, or a lower quality bilingual dictionary, BISPARSE-DEP degrades only marginally. In all these cases, it compares favorably with models that have been supplied with all necessary resources, showing promise for low-resource settings. We extensively evaluate BISPARSE-DEP on a new crowd-sourced cross-lingual dataset, with over 2900 hypernym pairs, spanning four languages from distinct families – French, Russian, Arabic and Chinese – and release the datasets for future evaluations.

2 Related Work

Cross-lingual Distributional Semantics

Cross-lingual word embeddings have been shown to encode semantics across languages in tasks such as word similarity (Faruqui and Dyer, 2014) and lexicon induction (Vulić and Moens, 2015). Our work stands apart in two aspects (1) In contrast to tasks involving similarity and synonymy (symmetric relations), the focus of our work is on detecting *asymmetric* relations across languages, using cross-lingual embeddings. (2) Unlike most previous work, we use dependency context instead of lexical context to induce cross-lingual embeddings, which allows us to abstract away from language specific word order, and (as we show) improves hypernymy detection.

More closely related is our prior work (Vyas and Carpuat, 2016) where we used lexical context based embeddings to detect cross-lingual lexical entailment. In contrast, the focus of this work is on hypernymy, a more well-defined relation than entailment. Also, we improve upon our previous approach by using dependency based embeddings (§6.1), and show that the improvements hold even when exposed to data scarce settings (§6.3).




Figure 2: Example Dependency Tree.

We also do a more comprehensive evaluation on four languages paired with English, instead of just French.

Dependency Based Embeddings In monolingual settings, dependency based embeddings have been shown to outperform window based embeddings on many tasks (Bansal et al., 2014; Hill et al., 2014; Melamud et al., 2016). Roller and Erk (2016) showed that dependency embeddings can help in recovering Hearst patterns (Hearst, 1992) like “animals such as cats”, which are known to be indicative of hypernymy. Shwartz et al. (2017) demonstrated that dependency based embeddings are almost always superior to window based embeddings for identifying hypernyms in English. Our work uses dependency based embeddings in a cross-lingual setting, a less explored research direction. A key novelty of our work also lies in its use of syntactic transfer to derive dependency contexts. This scenario is more relevant in a cross-lingual setting, where treebanks might not be available for many languages.

3 Our Approach – BISPARSE-DEP

We propose BISPARSE-DEP, a family of approaches that uses sparse, bilingual, dependency based word embeddings to identify cross-lingual hypernymy.

Figure 1 shows an overview of the end-to-end pipeline of BISPARSE-DEP. The two key components of this pipeline are: (1) *Dependency based contexts* (§3.1), which help us generalize across languages with minimal customization by abstracting away language-specific word order. We also discuss how to extract such contexts in the absence of a treebank in the language (§3.2) using a (weak) dependency parser trained on related languages. (2) *Bilingual sparse coding* (§3.3), which allows us to align dependency based word embeddings in a shared semantic space using a small bilingual dictionary. The resulting sparse bilingual embeddings can then be used with an unsupervised entailment scorer (§3.4) to predict hypernymy for cross-lingual word pairs.

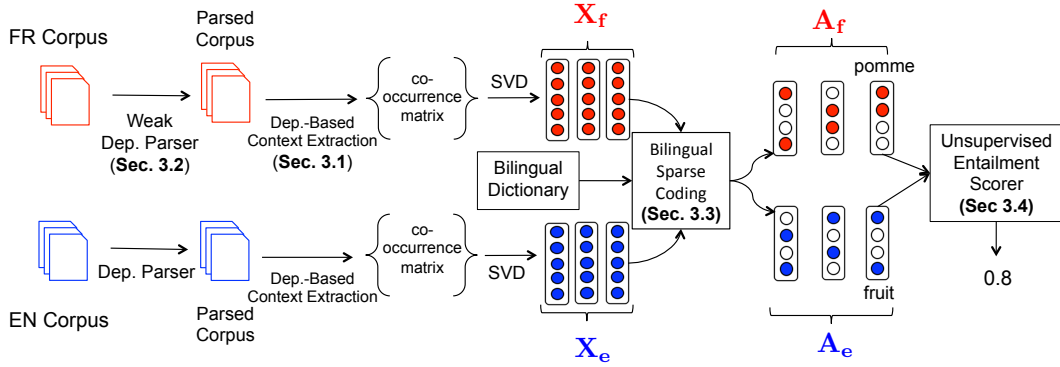


Figure 1: The BiSPARSE-DEP approach, which learns sparse bilingual embeddings using dependency based contexts. The resulting sparse embeddings, together with an unsupervised entailment scorer, can detect hypernyms across languages (e.g., *pomme* is a *fruit*).

3.1 Dependency Based Context Extraction

The context of a word can be described in multiple ways using its syntactic neighborhood in a dependency graph. For instance, in Figure 2, we describe the context for a target word (*traveler*) in the following two ways:

- **FULL context** (Padó and Lapata, 2007; Baroni and Lenci, 2010; Levy and Goldberg, 2014): Children and parent words, concatenated with the label and direction of the relation (eg. *roamed#nsubj*⁻¹ and *tired#amod* are contexts for *traveler*).
- **JOINT context** (Chersoni et al., 2016): Parent concatenated with each of its siblings (eg. *roamed#desert* and *roamed#seeking* are contexts for *traveler*).

These two contexts exploit different amounts of syntactic information – JOINT does not require labeled parses, unlike FULL. The JOINT context combines parent and sibling information, while FULL keeps them as distinct contexts. Both encode directionality into the context, either through label direction or through sibling-parent relations.

We use word-context co-occurrences generated using these contexts in a distributional semantic model (DSM) in lieu of window based contexts to generate dependency based embeddings.

3.2 Dependency Contexts without a Treebank

Using dependency contexts in multilingual settings may not always be possible, as dependency treebanks are not available for many languages. To circumvent this issue, we use related languages to train a weak dependency parser.

We train a *delexicalized* parser using treebanks of related languages, where the word form based

features are turned off, so that the parser is trained on purely non-lexical features (e.g. POS tags). The rationale behind this is that related languages show common syntactic structure that can be transferred to the original language, with delexicalized parsing (Zeman and Resnik, 2008; McDonald et al., 2011, *inter alia*) being one popular approach.³

3.3 Bilingual Sparse Coding

Given a dependency based co-occurrence matrix described in the previous section(s), we generate BiSPARSE-DEP embeddings using the framework from our prior work (Vyas and Carpuat, 2016), which we henceforth call BiSPARSE. BiSPARSE generates sparse, bilingual word embeddings using a dictionary learning objective with a sparsity inducing l_1 penalty. We give a brief overview of this approach, the full details of which can be found in our prior work.

For two languages with vocabularies v_e and v_f , and monolingual dependency embeddings \mathbf{X}_e and \mathbf{X}_f , BiSPARSE solves the following objective:

$$\begin{aligned}
 \operatorname{argmin}_{\mathbf{A}_e, \mathbf{D}_e, \mathbf{A}_f, \mathbf{D}_f} & \sum_{i=1}^{v_e} \frac{1}{2} \|\mathbf{A}_{e_i} \mathbf{D}_e^T - \mathbf{X}_{e_i}\|_2^2 + \lambda_e \|\mathbf{A}_{e_i}\|_1 \\
 & + \sum_{j=1}^{v_f} \frac{1}{2} \|\mathbf{A}_{f_j} \mathbf{D}_f^T - \mathbf{X}_{f_j}\|_2^2 + \lambda_f \|\mathbf{A}_{f_j}\|_1 \\
 & + \sum_{i,j} \frac{1}{2} \lambda_x \mathbf{S}_{ij} \|\mathbf{A}_{e_i} - \mathbf{A}_{f_j}\|_2^2 \quad (1) \\
 \text{s.t. } & \mathbf{A}_k > \mathbf{0} \quad \|\mathbf{D}_{k_i}\|_2^2 \leq 1 \quad \mathbf{k} \in \{e, f\}
 \end{aligned}$$

where \mathbf{S} is a translation matrix, and \mathbf{A}_e and \mathbf{A}_f

³More sophisticated techniques for transferring syntactic knowledge have been proposed (Ammar et al., 2016; Rasooli and Collins, 2017), but we prioritize simplicity and show that a simple delexicalized parser is effective.

are sparse matrices which are bilingual representations in a shared semantic space. The translation matrix \mathbf{S} (of size $v_e \times v_f$) captures correspondences between the vocabularies (of size v_e and v_f) of two languages. For instance, each row of \mathbf{S} can be a one-hot vector that identifies the word in f that is most frequently aligned with the e word for that row in a large parallel corpus, thus building a one-to-many mapping between the two languages.

3.4 Unsupervised Entailment Scorer

A variety of scorers can be used to quantify the directional relationship between two words, given feature representations of these words (Lin, 1998; Weeds and Weir, 2003; Lenci and Benotto, 2012). Once the BiSPARSE-DEP embeddings are constructed, we use *BalAPinc* (Kotlerman et al., 2009) to score word pairs for hypernymy. *BalAPinc* is based on the distributional inclusion hypothesis (Geffet and Dagan, 2005) and computes the geometric mean of 1) *LIN* (Lin, 1998), a symmetric score that captures similarity, and 2) *APinc*, an asymmetric score based on average precision.

4 Crowd-Sourcing Annotations

There is no publicly available dataset to evaluate models of hypernymy detection across multiple languages. While ontologies like Open Multilingual WordNet (OMW) (Bond and Foster, 2013) and BabelNet (Navigli and Ponzetto, 2012) contain cross-lingual links, these resources are semi-automatically generated and hence contain noisy edges. Thus, to get reliable and high-quality test beds, we collect evaluation datasets using CrowdFlower⁴. Our datasets span four languages from distinct families - French (Fr), Russian (Ru), Arabic (Ar) and Chinese (Zh) - paired with English.

To begin the annotation process, we first pool candidate pairs using hypernymy edges across languages from OMW and BabelNet, along with translations from monolingual hypernymy datasets (Baroni and Lenci, 2011; Baroni et al., 2012; Kotlerman et al., 2010).

4.1 Annotation Setup

The annotation task requires annotators to be fluent in both English and the non-English language. To ensure only fluent speakers perform the task, for each language, we provide task instructions in the non-English language itself. Also, we restrict the task to annotators verified by CrowdFlower to have those language skills. Finally, annotators also

⁴<http://crowdflower.com>

pair	#crowdsourced	#pos (= #neg)
French-English	2115	763
Russian-English	2264	706
Arabic-English	2144	691
Chinese-English	2165	806

Table 1: Crowd-sourced dataset statistics. #pos (#neg) denote positives (negatives) in the evaluation set. We deliberately under-sample negatives to have a balanced evaluation set.

need to pass a quiz based on a small amount of gold standard data to gain access to the task.

Annotators choose between three options for each word pair (p_f, q_e) , where p_f is a non-English word and q_e is a English word : “ p_f is a kind of q_e ”, “ q_e is a part of p_f ” and “none of the above”. Word pairs labeled with the first option are considered as positive examples while those labeled as “none of the above” are considered as negative.⁵ The second option was included to filter out meronymy examples that were part of the noisy pool. We leave it to the annotator to infer whether the relation holds between any senses of p_f or q_e , if either of them are polysemous.

For every candidate hypernym pair (p_f, q_e) , we also ask annotators to judge its reversed and translated *hyponym* pair (q_f, p_e) . For instance, if $(citron, food)$ is a hypernym candidate, we also show annotators $(aliments, lemon)$ which is a potential hyponym candidate (*potential*, because as mentioned in §1, translation need not preserve semantic relationships). The purpose of presenting the hyponym pair, (q_f, p_e) , is two-fold. First, it emphasizes the directional nature of the task. Second, it identifies hyponym pairs, which we use as negative examples. The hyponym pairs are challenging since differentiating them from hypernyms truly requires detecting asymmetry.

Each pair was judged by at least 5 annotators, and judgments with 80% agreement (at least 4 annotators agree) are considered for the final dataset. This is a stricter condition than certain monolingual hypernymy datasets - for instance, EVALution (Santus et al., 2015) - where agreement by 3 annotators is deemed sufficient. Inter-annotator agreement measured using Fleiss’ Kappa (Fleiss, 1971) was 58.1 (French), 53.7 (Russian), 53.2 (Arabic) and 55.8 (Chinese). This indicates moderate agreement, on par with agreement obtained on related fine-grained semantic tasks (Pavlick et al., 2015). We cannot compare with monolin-

⁵We collected more negative pairs than positive, but sampled so as to keep a balanced dataset for ease of evaluation. We will release all annotated pairs along with the dataset.

gual hypernymy annotator agreement as, to the best of our knowledge, such numbers are not available for existing test sets. Dataset statistics are shown in Table 1.

We observed that annotators were able to agree on pairs containing polysemous words where hypernymy holds for some sense. For instance, for the French-English pair (*avocat, professional*), the French word *avocat* can either mean *lawyer* or *avocado*, but the pair was annotated as a positive example. Hence, we leave it to the annotators to handle polysemy by choosing the most appropriate sense.

4.2 Two Evaluation Test Sets

To verify if the crowdsourced hyponyms are challenging negative examples we create two evaluation sets. Both share the (crowdsourced) positive examples, but differ in their negatives:

- HYPER-HYPO – negative examples are the crowdsourced hyponyms.
- HYPER-COHYPO – negative examples are *cohyponyms* drawn from OMW.

Cohyponyms are words sharing a common hypernym. For instance, *bière* (“beer” in French) and *vodka* are cohyponyms since they share a common hypernym in *alcoollalcohol*. We choose cohyponyms for the second test set because: (a) They require differentiating between similarity (a symmetric relation) and hypernymy (an asymmetric relation). For instance, *bière* and *vodka* are highly similar yet, they do not have a hypernymy relationship. (b) Cohyponyms are a popular choice of negative examples in many entailment datasets (Baroni and Lenci, 2011).

5 Experimental Setup

5.1 Data and Evaluation Setup

Training BISPARE-DEP requires a dependency parsed monolingual corpus, and a translation matrix for jointly aligning the monolingual vectors. We compute the translation matrix using word alignments derived from parallel corpora (see corpus statistics in Table ??). While we use parallel corpora to generate the translation matrix to be comparable to baselines (§5.2), we can obtain the matrix from any bilingual dictionary.

The monolingual corpora are parsed using Yara Parser (Rasooli and Tetreault, 2015), trained on the corresponding treebank from the Universal Dependency Treebank (McDonald et al., 2013) (UDT-v1.4). Yara Parser was

chosen as it is fast, and competitive with state-of-the-art parsers (Choi et al., 2015). The monolingual corpora was POS-tagged using TurboTagger (Martins et al., 2013). We induce dependency contexts for words by first thresholding the language vocabulary to the top 50,000 nouns, verbs and adjectives. A co-occurrence matrix is computed over this vocabulary using the context types in §3.1.

Inducing Dependency Contexts The entries of the word-context co-occurrence matrix are re-weighted using Positive Pointwise Mutual Information (Bullinaria and Levy, 2007). The resulting matrix is reduced to 1000 dimensions using SVD (Golub and Kahan, 1965).⁶ These vectors are used as \mathbf{X}_e , \mathbf{X}_f in the setup from §3.3 to generate 100 dimensional sparse bilingual vectors.

Evaluation We use accuracy as our evaluation metric, as it is easy to interpret when the classes are balanced (Turney and Mohammad, 2015). Both evaluation datasets – HYPER-HYPO and HYPER-COHYPO – are split into 1:2 dev/test splits. *BalAPinc* has two tunable parameters - 1) a threshold that indicates the *BalAPinc* score above which all examples are labeled as positive, 2) the maximum number of features to consider for each word. We use the tuning set to tune the two parameters as well as the various hyper-parameters associated with the models.

5.2 Contrastive Approaches

We compare our BISPARE-DEP embeddings with the following approaches:

MONO-DEP (Translation baseline) For word pair (p_f , q_e) in test data, we translate p_f to English using the most common translation in the translation matrix. Hypernymy is then determined using sparse, dependency based embeddings in English.

BISPARE-LEX (Window context) Predecessor of the BISPARE-DEP model from our previous work (Vyas and Carpuat, 2016). This model induces sparse, cross-lingual embeddings using window based context.

BIVEC+ (Window context) Our extension of the BIVEC model of Luong et al. (2015). BIVEC generates dense, cross-lingual embeddings using window based context, by substituting aligned word pairs within a window in parallel sentences. By default, BIVEC only trains using parallel data,

⁶Chosen based on preliminary experiments with {500,1000,2000,3000} dimensional vectors for En-Fr.

Language	Parallel Data	#sent.	Monolingual Data	#sent.
English	–	–	Wackypedia (Baroni et al., 2009)	43M
Arabic	ISI (Munteanu and Marcu, 2007) NewsCommentary, Wikipedia (Tiedemann, 2012)	1.1M	Arabic Gigaword 3.0 (Graff, 2007)	17M
Chinese	FBIS (LDC2003E14)	9.5M	Chinese Gigaword 5.0 (Parker, 2011)	58M
French	Europarl (Koehn, 2005) NewsCommentary [♦] , Wikipedia (Tiedemann, 2012)	2.7M	Wikipedia [♣]	20M
Russian	Yandex-1M [♠]	1.6M	Wikipedia [♣]	22M

♦ = www.statmt.org/wmt15/training-parallel-nc-v10.tgz, ♣ = dumps.wikimedia.org/xxwiki/20161201/
♠ = translate.yandex.ru/corpus

Table 2: Training data statistics for different languages. Note that while we use parallel corpora for computing translation dictionaries, our approach does not require it, and can work with any bilingual dictionary.

and so we initialize it with monolingually trained window based embeddings to ensure fair comparison.

CL-DEP (Dependency context) The model from Vulić (2017), which induces dense, dependency based cross-lingual embeddings by translating syntactic word-context pairs using the most common translation, and jointly training a `word2vecf`⁷ model for both languages. Vulić (2017) showed improvements for word similarity and bilingual lexicon induction. We report the first results using CL-DEP on this task.

5.3 Evaluating Robustness of BISPARSE-DEP

We investigate how robust BISPARSE-DEP is when exposed to data scarce settings. Evaluating on a truly low resource language is complicated by the fact that obtaining an evaluation dataset for such a language is difficult. Therefore, we simulate such settings for the languages in our dataset in multiple ways.

No Treebank If a treebank is not available for a language, dependency contexts have to be induced using treebanks from other languages (§3.2), which can affect the quality of the dependency-based embeddings. To simulate this, we train a delexicalized parser for the languages in our dataset. We use treebanks from Slovenian, Ukrainian, Serbian, Polish, Bulgarian, Slovak and Czech (40k sentences) for training the Russian parser, and treebanks from English, Spanish, German, Portuguese, Swedish and Italian (66k sentences) for training the French parser. UDT does not (yet) have languages in the same family as Arabic or Chinese, so for the sake of completeness, we train Arabic and Chinese parsers on delexicalized treebanks of the language itself. Af-

⁷bitbucket.org/yoavgo/word2vecf/

ter delexicalized training, the Labeled Attachment Score (LAS) on the UDT test set dropped by several points for all languages – from 76.6% to 60.0% for Russian, 83.7% to 71.1% for French, from 76.3% to 62.4% for Arabic and from 80.3% to 53.3% for Chinese. The monolingual corpora are then parsed with these weaker parsers, and co-occurrences and dependency contexts are computed as before.

Subsampling Monolingual Data To simulate low-resource behavior along another axis, we subsample the monolingual corpora used by BISPARSE-DEP to induce monolingual vectors, \mathbf{X}_e , \mathbf{X}_f . Specifically, we learn \mathbf{X}_e and \mathbf{X}_f using progressively smaller corpora.

Quality of Bilingual Dictionary We study the impact of the quality of the bilingual dictionary used to create the translation matrix \mathbf{S} . This experiment involves using increasingly smaller parallel corpora to induce the translation dictionary.

6 Experiments

We aim to answer the following questions – (a) Are dependency based embeddings superior to window based embeddings for identifying cross-lingual hypernymy? (§6.1) (b) Does directionality in the dependency context help cross-lingual hypernymy identification? (§6.2) (c) Are our models robust in data scarce settings (§6.3)? (d) Is the answer to (a) predicated on the choice of entailment scorer? (§6.4)?

6.1 Dependency v/s Window Contexts

We compare the performance of models described in §5.2 with the BISPARSE-DEP (FULL and JOINT) models. We evaluate the models on the two test splits described in §4.2 – HYPER-HYPO and HYPER-COHYPO.

Model ↓ En With →	Ru	Zh	Ar	Fr	Avg.
Translation Baseline					
MONO-DEP	50.1	52.3	51.8	54.5	52.2
Win. Based					
BiSPARSE-LEX	56.6	53.7	50.9	52.0	53.3
BIVEC+	55.8	52.0	51.5	53.4	53.2
Dep. Based					
CL-DEP	60.2	54.4	56.7*	53.8	56.3
BiSPARSE-DEP (Full)	59.0	55.9	52.6	56.6	56.0
BiSPARSE-DEP (Joint)	53.8	57.0*	52.4	59.9*	55.8
BiSPARSE-DEP (Unlab)	55.9	51.2	53.3	55.9	54.1

(a) Performance on HYPER-HYPO.

Model ↓ En With →	Ru	Zh	Ar	Fr	Avg.
Translation Baseline					
MONO-DEP	58.7	50.0	65.1	56.9	57.7
Win. Based					
BiSPARSE-LEX	63.8	55.8	65.8	63.2	62.2
BIVEC+	55.9	64.9	62.2	54.1	58.3
Dep. Based					
CL-DEP	56.2	62.7	63.1	61.0	60.0
BiSPARSE-DEP (Full)	63.6	67.3	66.8*	66.7*	66.1
BiSPARSE-DEP (Joint)	60.6	63.6	65.9	64.9	63.8
BiSPARSE-DEP (Unlab)	58.6	66.7	62.4	61.5	62.4

(b) Performance on HYPER-COHYPO.

Table 3: Comparing the different approaches from §5.2 with our BiSPARSE-DEP approach on HYPER-HYPO and HYPER-COHYPO (random baseline= 0.5). **Bold** denotes the best score for each language, and the * on the best score indicates a statistically significant ($p < 0.05$) improvement over the next best score, using McNemar’s test (McNemar, 1947). Across both datasets, BiSPARSE-DEP models outperform window based models and the translation baseline on an average.

Hyper-Hypo Results Table 3a shows the results on HYPER-HYPO. First, the benefit of cross-lingual modeling (as opposed to translation) is evident in that almost all models (except CL-DEP on French) outperform the translation baseline. Among dependency based models, BiSPARSE-DEP (FULL) and CL-DEP consistently outperform both window models, while BiSPARSE-DEP (JOINT) outperforms them on all except Russian. BiSPARSE-DEP (JOINT) is the best model overall for two languages (French and Chinese), CL-DEP for one (Arabic), with no statistically significant differences between BiSPARSE-DEP (JOINT) and CL-DEP for Russian. This confirms that dependency context is more useful than window context for cross-lingual hypernymy detection.

Hyper-Cohypo Results The trends observed on HYPER-HYPO also hold on HYPER-COHYPO i.e. dependency based models continue to outperform window based models (Table 3b).

Overall, BiSPARSE-DEP (FULL) performs best in this setting, followed closely by BiSPARSE-DEP (JOINT). This suggests that the sibling information encoded in JOINT is useful to distinguish hypernyms from hyponyms (HYPER-HYPO results), while the dependency labels encoded in FULL help to distinguish hypernyms from cohyponyms. Also note that all models improve significantly on the HYPER-COHYPO set, suggesting that discriminating hypernyms from cohyponyms is easier than discriminating them from hyponyms.

While the BiSPARSE-DEP models were generally performing better than window models on both test sets, CL-DEP was not as consistent (e.g.,

it was worse than the best window model on HYPER-COHYPO). As shown by Turney and Mohammad (2015), *BalAPinc* is designed for sparse embeddings and is likely to perform poorly with dense embeddings. This explains the relatively inconsistent performance of CL-DEP.

Besides establishing the challenging nature of our crowd-sourced set, the experiments on HYPER-COHYPO and HYPER-HYPO also demonstrate the ability of the BiSPARSE-DEP models to discriminate between different lexical semantic relations (viz. hypernymy and cohyponymy) in a cross-lingual setting. We will investigate this ability more carefully in future work.

6.2 Ablating Directionality in Context

The context described by the FULL and JOINT BiSPARSE models encodes directional information (§3.1) either in the form of label direction (FULL), or using sibling information (JOINT). Does such directionality in the context help to capture the asymmetric relationship inherent to hypernymy? To answer this, we evaluate a third BiSPARSE-DEP model which uses UNLABELED dependency contexts. This is similar to the FULL context, except we do not concatenate the label of the relation to the context word (parent or children). For instance, for *traveler* in Fig. 2, contexts will be *roamed* and *tired*.

Experiments on both HYPER-HYPO and HYPER-COHYPO (bottom row, Tables 3a and 3b) highlight that directional information is indeed essential - UNLABELED almost always performs worse than FULL and JOINT, and in many cases worse than even window based models.

Model ↓ En With →	Ru	Zh	Ar	Fr	Avg.
Hyper-Hypo					
Best Win.	56.6	53.7	51.5	53.4	53.8
Delex.	59.1*	55.1*	54.6*	56.1*	56.2
Best Dep.	60.2	57.0*	56.7*	59.9*	58.5
Hyper-Cohypo					
Best Win.	63.8	64.9	65.8	63.2	64.4
Delex.	59.4	65.7*	67.5*	66.3*	64.7
Best Dep.	63.6*	67.3*	66.8*	66.7	66.1

Table 4: **Robustness in absence of a treebank:** The delexicalized model is competitive to the best dependency based and the best window based models on both test sets. For each dataset, * indicates a statistically significant ($p < 0.05$) improvement over the next best model in that column, using McNemar’s test (McNemar, 1947).

6.3 Evaluating Robustness of BISPARSE-DEP

No Treebank We run experiments (Table 4) for all languages with a version of BISPARSE-DEP that use the FULL context type for both English and the non-English (target) language, but the target language contexts are derived from a corpus parsed using a delexicalized parser (§5.3).

This model compares favorably on all language pairs against the best window based and the best dependency based model. In fact, it almost consistently outperforms the best window based model by several points, and is only slightly worse than the best dependency-based model.

Further analysis revealed that the good performance of the delexicalized model is due to the relative robustness of the delexicalized parser on frequent contexts in the co-occurrence matrix. Specifically, we found that in French and Russian, the most frequent contexts were derived from *amod*, *nmod*, *nsubj* and *doobj* edges.⁸ For instance, the *nmod* edge appears in 44% of Russian contexts and 33% of the French contexts. The delexicalized parser predicts both the label and direction of the *nmod* edge correctly with an F1 of 68.6 for Russian and 69.6 for French. In contrast, a fully-trained parser achieves a F1 of 76.7 for Russian and 76.8 for French for the same edge.

Small Monolingual Corpus In Figure 4, we use increasingly smaller monolingual corpora (10%, 20%, 40%, 60% and 80%) sampled at random to induce the monolingual vectors for BISPARSE-DEP (FULL) model. Trends (Figure 4) indicate that BISPARSE-DEP models that use only 40% of the original data remain competitive with the BISPARSE-LEX model that has access to the full

⁸Together they make up at least 70% of the contexts.




Figure 3: **Robustness to Small Corpus** For most languages, BISPARSE-DEP outperforms the corresponding best window based model for each language on HYPER-HYPO, with about 40% of the monolingual corpora.




Figure 4: **Robustness to Noisy Dictionary** For most languages, BISPARSE-DEP outperforms the corresponding best window based model on HYPER-HYPO, with increasingly lower quality dictionaries.

data. Robust performance with smaller monolingual corpora is helpful since large-enough monolingual corpora are not always easily available.

Quality of Bilingual Dictionary Bilingual dictionaries derived from smaller amounts of parallel data are likely to be of lower quality than those derived from larger corpora. Hence, to analyze the impact of dictionary quality on BISPARSE-DEP (FULL), we use increasingly smaller parallel corpora to induce bilingual dictionaries used as the score matrix S (§3.3). We use the top 10%, 20%, 40%, 60% and 80% sentences from the parallel corpora. The trends in Figure 4 show that even with a lower quality dictionary, BISPARSE-DEP performs better than BISPARSE-LEX.

6.4 Choice of Entailment Scorer

We change the entailment scorer from *BalAPinc* to *SLQS* (Santus et al., 2014) and redo experiments from §6.1 to see if the conclusions drawn depend

on the choice of the entailment scorer. *SLQS* is based on the distributional informativeness hypothesis, which states that hypernyms are less “informative” than hyponyms, because they occur in more general contexts. The informativeness E_u of a word u is defined to be the median entropy of its top N dimensions, $E_u = \text{median}_{k=1}^N H(c_k)$, where $H(c_i)$ denotes the entropy of dimension c_i . The *SLQS* score for a pair (u, v) is the relative difference in entropies,

$$SLQS(u \rightarrow v) = 1 - \frac{E_u}{E_v}$$

Recent work (Shwartz et al., 2017) has found *SLQS* to be more successful than other metrics in monolingual hypernymy detection.

The trends observed in these experiments are consistent with those in §6.1 – both BISPARSE-DEP models still outperform window-based models. Also, the delexicalized version of BISPARSE-DEP outperforms the window-based models, showing that the robust behavior demonstrated in §6.3 is also invariant across metrics.

We also found that using *BalAPinc* led to better results than *SLQS*. For both BISPARSE-DEP models, *BalAPinc* wins across the board for two languages (Russian and Chinese), and wins half the time for the other two languages compared to *SLQS*. We leave detailed comparison of these and other scores to future work.

7 Conclusion

We introduced BISPARSE-DEP, a new distributional approach for identifying cross-lingual hypernymy, based on cross-lingual embeddings derived from dependency contexts. We showed that using BISPARSE-DEP is superior for the cross-lingual hypernymy detection task, when compared to standard window based models and a translation baseline. Further analysis also showed that BISPARSE-DEP is robust to various low-resource settings. In principle, BISPARSE-DEP can be used for any language that has a bilingual dictionary with English and a “related” language with a treebank. We also introduced crowd-sourced cross-lingual hypernymy datasets for four languages for future evaluations.

Our approach has the potential to complement existing work on creating cross-lingual ontologies such as BabelNet and the Open Multilingual Wordnet, which are noisy because they are compiled semi-automatically, and have limited language coverage. In general, distributional approaches can help refine ontology construction for

any language where sufficient resources are available.

It remains to be seen how our approach performs for other language pairs beyond simulated low-resource settings. We anticipate that replacing our delexicalized parser with more sophisticated transfer strategies (Rasooli and Collins, 2017; Aufrant et al., 2016) might be beneficial in such settings. While our delexicalized parsing based approach exhibits robustness, it can benefit from more sophisticated approaches for transfer parsing (Rasooli and Collins, 2017; Aufrant et al., 2016) to improve parser performance. We aim to explore these and other directions in the future.

Acknowledgments

The authors would like to thank the members of the CLIP lab at the University of Maryland, members of the Cognitive Computation Group at the University of Pennsylvania, and the anonymous reviewers from EMNLP/CoNLL 2017 and NAACL 2018 for their constructive feedback. YV and MC were funded in part by research awards from Amazon, Google, and the Clare Boothe Luce Foundation. SU and DR were supported by Contract HR0011-15-2-0025 with the US Defense Advanced Research Projects Agency (DARPA).

References

- Waleed Ammar, George Mulcaire, Miguel Ballesteros, Chris Dyer, and Noah Smith. 2016. Many languages, one parser. *Transactions of the Association for Computational Linguistics* 4:431–444.
- Lauriane Aufrant, Guillaume Wisniewski, and François Yvon. 2016. Zero-resource dependency parsing: Boosting delexicalized cross-lingual transfer with linguistic knowledge. In *Proc. of COLING*. The COLING 2016 Organizing Committee, Osaka, Japan, pages 119–130. <http://aclweb.org/anthology/C16-1012>.
- Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2014. Tailoring continuous word representations for dependency parsing. In *Proc. of ACL*. Association for Computational Linguistics, Baltimore, Maryland, pages 809–815. <http://www.aclweb.org/anthology/P14-2131>.
- Marco Baroni, Raffaella Bernardi, Ngoc-Quynh Do, and Chung-chieh Shan. 2012. Entailment above the word level in distributional semantics. In *Proc. of EACL*. Association for Computational Linguistics, pages 23–32. <http://www.aclweb.org/anthology/E12-1004>.
- Marco Baroni, Silvia Bernardini, Adriano Ferraresi, and Eros Zanchetta. 2009. The WaCky wide web:

- a collection of very large linguistically processed web-crawled corpora. *Proc. of LREC* .
- 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. How we BLESSed distributional semantic evaluation. In *Proceedings of the GEMS 2011 Workshop*.
- Francis Bond and Ryan Foster. 2013. Linking and extending an open multilingual wordnet. In *Proc. of ACL*. Association for Computational Linguistics, pages 1352–1362. <http://www.aclweb.org/anthology/P13-1133>.
- John A. Bullinaria and Joseph P. Levy. 2007. Extracting semantic representations from word co-occurrence statistics: A computational study. *Behavior Research Methods* pages 510–526.
- Vicente Ivan Sanchez Carmona and Sebastian Riedel. 2017. How well can we predict hypernyms from word embeddings? a dataset-centric analysis. In *Proc. of EACL*. Association for Computational Linguistics, Valencia, Spain, pages 401–407. <http://www.aclweb.org/anthology/E17-2064>.
- Emmanuele Chersoni, Enrico Santus, Alessandro Lenci, Philippe Blache, and Chu-Ren Huang. 2016. Representing verbs with rich contexts: an evaluation on verb similarity. In *Proc. of EMNLP*. Association for Computational Linguistics, Austin, Texas, pages 1967–1972. <https://aclweb.org/anthology/D16-1205>.
- Jinho D Choi, Joel R Tetreault, and Amanda Stent. 2015. It depends: Dependency parser comparison using a web-based evaluation tool. In *Proc. of ACL*. Association for Computational Linguistics, Beijing, China, pages 387–396. <http://www.aclweb.org/anthology/P15-1038>.
- Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proc. of EACL*. Association for Computational Linguistics, Gothenburg, Sweden, pages 462–471. <http://www.aclweb.org/anthology/E14-1049>.
- Joseph L Fleiss. 1971. Measuring nominal scale agreement among many raters. *Psychological bulletin* 76(5):378.
- Ruiji Fu, Jiang Guo, Bing Qin, Wanxiang Che, Haifeng Wang, and Ting Liu. 2014. Learning semantic hierarchies via word embeddings. In *Proc. of ACL*. Association for Computational Linguistics, Baltimore, Maryland, pages 1199–1209. <http://www.aclweb.org/anthology/P14-1113>.
- Maayan Geffet and Ido Dagan. 2005. The Distributional Inclusion Hypotheses and Lexical Entailment. In *Proc. of ACL*.
- Tom Goldstein, Christoph Studer, and Richard Baraniuk. 2014. A Field Guide to Forward-Backward Splitting with a FASTA Implementation. *arXiv eprint* abs/1411.3.
- Gene Golub and William Kahan. 1965. Calculating the singular values and pseudo-inverse of a matrix. *Journal of the SIAM* .
- David Graff. 2007. Arabic gigaword 3rd edition, LDC2003T40. LDC, University of Pennsylvania.
- Marti A. Hearst. 1992. Automatic acquisition of hyponyms from large text corpora. In *Proc. of COLING*. Association for Computational Linguistics, pages 539–545. <https://doi.org/10.3115/992133.992154>.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2014. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *arXiv preprint arXiv:1408.3456* .
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *Proc. of MT Summit*.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2009. Directional distributional similarity for lexical expansion. In *Proc. of the ACL-IJCNLP*.
- Lili Kotlerman, Ido Dagan, Idan Szpektor, and Maayan Zhitomirsky-Geffet. 2010. Directional distributional similarity for lexical inference. *Natural Language Engineering* .
- Alessandro Lenci and Giulia Benotto. 2012. Identifying hypernyms in distributional semantic spaces. In *Proc. of the 6th Workshop on Semantic Evaluation*.
- Omer Levy and Yoav Goldberg. 2014. Dependency-based word embeddings. In *Proc. of ACL*. Association for Computational Linguistics, Baltimore, Maryland, pages 302–308. <http://www.aclweb.org/anthology/P14-2050>.
- Dekang Lin. 1998. Automatic retrieval and clustering of similar words. In *Proc. of ACL*. Association for Computational Linguistics, pages 768–774. <https://doi.org/10.3115/980691.980696>.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Bilingual word representations with monolingual quality in mind. In *Proc. of the Workshop on Vector Space Modeling for NLP*.
- Andre Martins, Miguel Almeida, and Noah A. Smith. 2013. Turning on the Turbo: Fast Third-Order Non-Projective Turbo Parsers. In *Proc. of ACL*. Association for Computational Linguistics, pages 617–622. <http://www.aclweb.org/anthology/P13-2109>.
- Ryan McDonald, Joakim Nivre, Yvonne Quirnbach-Brundage, Yoav Goldberg, Dipanjan Das, Kuzman Ganchev, Keith Hall, Slav Petrov, Hao

- Zhang, Oscar Täckström, Claudia Bedini, Núria Bertomeu Castelló, and Jungmee Lee. 2013. Universal dependency annotation for multilingual parsing. In *Proc. of ACL*.
- Ryan McDonald, Slav Petrov, and Keith Hall. 2011. Multi-source transfer of delexicalized dependency parsers. In *Proc. of EMNLP*. Association for Computational Linguistics, pages 62–72. <http://www.aclweb.org/anthology/D11-1006>.
- Quinn McNemar. 1947. Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika* 12(2):153–157.
- Oren Melamud, David McClosky, Siddharth Patwardhan, and Mohit Bansal. 2016. The Role of Context Types and Dimensionality in Learning Word Embeddings. In *Proc. of NAACL-HLT*. Association for Computational Linguistics, San Diego, California, pages 1030–1040. <http://www.aclweb.org/anthology/N16-1118>.
- Dragos Stefan Munteanu and Daniel Marcu. 2007. ISI Arabic-English Automatically Extracted Parallel Text LDC2007T08. LDC, University of Pennsylvania.
- Roberto Navigli and Simone Paolo Ponzetto. 2012. BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *Artificial Intelligence*.
- Matteo Negri, Alessandro Marchetti, Yashar Mehdad, Luisa Bentivogli, and Danilo Giampiccolo. 2012. Semeval-2012 Task 8: Cross-lingual Textual Entailment for Content Synchronization.
- Matteo Negri, Alessandro Marchetti, Yashar Mehdad, Luisa Bentivogli, and Danilo Giampiccolo. 2013. Semeval-2013 Task 8: Cross-lingual Textual Entailment for Content Synchronization.
- Sebastian Padó, Michel Galley, Dan Jurafsky, and Chris Manning. 2009. Robust machine translation evaluation with entailment features. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP: Volume 1-Volume 1*. Association for Computational Linguistics, pages 297–305.
- Sebastian Padó and Mirella Lapata. 2007. Dependency-based construction of semantic space models. *Computational Linguistics*.
- Robert Parker. 2011. Chinese Gigaword 5th Edition, LDC2011T13. LDC, University of Pennsylvania.
- Ellie Pavlick, Pushpendre Rastogi, Juri Ganitkevitch, Benjamin Van Durme, and Chris Callison-Burch. 2015. PPDB 2.0: Better paraphrase ranking, fine-grained entailment relations, word embeddings, and style classification. *Proceedings of ACL-IJCNLP 2015* pages 425–430.
- Mohammad Sadegh Rasooli and Michael Collins. 2017. Cross-lingual syntactic transfer with limited resources. *Transactions of the Association for Computational Linguistics* 5:279–293. <https://transacl.org/ojs/index.php/tacl/article/view/922>.
- Mohammad Sadegh Rasooli and Joel R. Tetreault. 2015. Yara Parser: A Fast and Accurate Dependency Parser. *CoRR* abs/1503.06733.
- Stephen Roller and Katrin Erk. 2016. Relations such as Hypernymy: Identifying and Exploiting Hearst Patterns in Distributional Vectors for Lexical Entailment. In *Proc. of EMNLP*. Association for Computational Linguistics, pages 2163–2172. <http://www.aclweb.org/anthology/D16-1234>.
- Enrico Santus, Alessandro Lenci, Qin Lu, and Sabine Schulte im Walde. 2014. Chasing hypernyms in vector spaces with entropy. In *Proc. of EACL*. Association for Computational Linguistics, Gothenburg, Sweden, pages 38–42. <http://www.aclweb.org/anthology/E14-4008>.
- 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 4th Workshop on Linked Data in Linguistics (LDL-2015)*. Association for Computational Linguistics, Beijing, China, pages 64–69. <http://www.aclweb.org/anthology/W15-4208>.
- Vered Shwartz, Enrico Santus, and Dominik Schlechtweg. 2017. Hypernyms under siege: Linguistically-motivated artillery for hypernymy detection. In *Proc. of EACL*. Association for Computational Linguistics, Valencia, Spain, pages 65–75. <http://www.aclweb.org/anthology/E17-1007>.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in OPUS. In *Proc. of LREC*.
- Peter D Turney and Saif M Mohammad. 2015. Experiments with three approaches to recognizing lexical entailment. *Natural Language Engineering*.
- Piek Vossen, Egoitz Laparra, Itziar Aldabe, and German Rigau. 2015. Interoperability of cross-lingual and cross-document event detection. In *Proc. of the 3rd Workshop on EVENTS at the NAACL-HLT*.
- Ivan Vulić. 2017. Cross-lingual syntactically informed distributed word representations. In *Proc. of EACL*. Association for Computational Linguistics, Valencia, Spain, pages 408–414. <http://www.aclweb.org/anthology/E17-2065>.
- Ivan Vulić and Marie-Francine Moens. 2015. Bilingual word embeddings from non-parallel document-aligned data applied to bilingual lexicon induction. In *Proc. of ACL*. Association for Computational Linguistics, Beijing, China, pages 719–725. <http://www.aclweb.org/anthology/P15-2118>.

- Yogarshi Vyas and Marine Carpuat. 2016. *Sparse Bilingual Word Representations for Cross-lingual Lexical Entailment*. Association for Computational Linguistics, San Diego, California, pages 1187–1197. <http://www.aclweb.org/anthology/N16-1142>.
- Julie Weeds and David Weir. 2003. *A general framework for distributional similarity*. In *Proc. of EMNLP*. <http://www.aclweb.org/anthology/W03-1011>.
- Daniel Zeman and Philip Resnik. 2008. *Cross-language parser adaptation between related languages*. In *IJCNLP*. <http://www.aclweb.org/anthology/I08-3008>.