Embedding Words and Senses Together
via Joint Knowledge-Enhanced Training

Massimiliano Mancini*, Jose Camacho-Collados*, Ignacio Iacobacci and **Roberto Navigli**

Department of Computer Science
Sapienza University of Rome
mancini@dis.uniroma1.it
{collados,iacobacci,navigli}@di.uniroma1.it

Abstract

Word embeddings based on neural networks are widely used in Natural Language Processing. However, despite their success in capturing semantic information from massive corpora, word embeddings still conflate different meanings of a word into a single vectorial representation and do not benefit from information available in lexical resources. We address this issue by proposing a new model that jointly learns word and sense embeddings and represents them in a unified vector space by exploiting large corpora and knowledge obtained from semantic networks. We evaluate the main features of our approach qualitatively and quantitatively in various tasks, highlighting the advantages of the proposed method with respect to state-of-the-art word- and sense-based models.

1 Introduction

Recently, approaches based on neural networks which embed words into low-dimensional vector spaces from text corpora (i.e., word embeddings) have become increasingly popular (Mikolov et al., 2013; Pennington et al., 2014). Word embeddings have proved to be beneficial in many Natural Language Processing tasks such as Machine Translation (Zou et al., 2013), syntactic parsing (Weiss et al., 2015) or Question Answering (Bordes et al., 2014). However, these representations are generally hampered by an important limitation: the inability to discriminate among different meanings of the same word.

Previous works have addressed this limitation by automatically clustering word senses, repre-

senting them into separate vectors (Schütze, 1998; Reisinger and Mooney, 2010; Huang et al., 2012; Neelakantan et al., 2014). However, these approaches learn solely based on statistics extracted from text corpora and do not exploit knowledge from semantic networks. Additionally, their induced senses are not easily mappable to lexical resources, which limits their application. Recent approaches utilized semantic networks to inject knowledge into existing word representations (Yu and Dredze, 2014; Tian et al., 2014; Faruqui et al., 2015; Goikoetxea et al., 2015). Despite the increased exploitation of knowledge, the obtained embeddings are still unable to capture the different senses of a word, thus having a limited representation capability. In order to deal with these two issues, a number of approaches have leveraged lexical resources to learn sense embeddings as a result of post-processing conventional word embeddings (Chen et al., 2014; Johansson and Pina, 2015; Jauhar et al., 2015; Rothe and Schütze, 2015; Camacho-Collados et al., 2016).

In this paper we propose SW2V (*Senses and Words to Vectors*), a new model which, instead, simultaneously learns embeddings for both words and senses as an emerging feature by exploiting knowledge from both text corpora and semantic networks in a joint training phase. Our model provides three additional key features: (1) both word and sense embeddings are represented in the same vector space, (2) it is flexible as it can be applied to different predictive models, and (3) it is scalable for large semantic networks and large amounts of text corpora.

2 Related work

Embedding words from large corpora into a low-dimensional vector space has been a popular task since the appearance of the probabilistic feed-

---

Authors marked with an asterisk (*) contributed equally.

forward neural network language model (Bengio et al., 2003) and later developments such as Word2Vec (Mikolov et al., 2013). However, little research has focused on exploiting knowledge resources to overcome the inherent ambiguity of word embeddings.

A recent approach overcame this limitation by applying an off-the-shelf disambiguation system on a corpus and then using Word2Vec to learn sense embeddings over the pre-disambiguated text (Iacobacci et al., 2015, SensEmbed). However, in their approach words are replaced by their intended senses and consequently their representation is exclusively based on word senses. Two approaches have proposed a model to build a shared space of words and senses. Chen et al. (2014) proposed a model for obtaining word and sense representations based on a first training of conventional word embeddings, a disambiguation step based on sense definitions, and a final training phase that uses the disambiguated text as input. Likewise, Rothe and Schütze (2015, Autoextend) aimed at building a shared space of word and sense embeddings based on two steps: a first training of word embeddings and a second training to produce sense and synset embeddings. These two approaches require several steps of training and make use of a relatively small resource like WordNet, which limits their coverage and applicability. Instead, we put forward a new model for a single joint training of word and sense embeddings exploiting knowledge from both text corpora and large semantic networks, resulting in a common semantic space of word and sense embeddings.

3 Joint training of words and senses

The goal of our approach is to obtain a shared vector space of words and senses. To this end, our model extends conventional word embedding models by integrating explicit knowledge into its architecture. While we will focus on the Continuous Bag Of Words (CBOW) architecture of Word2Vec (Mikolov et al., 2013), our extension could be easily applied to Skip-Gram, or to other predictive approaches based on neural networks. The CBOW architecture is based on the feedforward neural network language model (Bengio et al., 2003) and aims at predicting the current word using its surrounding context. The architecture consists of input, hidden and output layers. The input layer has the size of the word vocabulary (V) and encodes the context as a combination of one-hot vector representations of surrounding words of a given target word. The output layer has the same size of the input layer and it is composed by a one-hot vector of the target word during the training phase.

Our model, SW2V, extends the input and output layers of the neural network with word senses[1] by exploiting the intrinsic connection between words and senses. In a given context, a word may be connected with zero, one or more senses. We will refer to the set of senses connected to a given word within the specific context as its *associated senses*. In Section 4 we explain how to obtain such connections.

Formally, we define a training instance as a sequence of words $W = w_{t-n}, ..., w_t, ..., w_{t+n}$ (being w_t the target word) and $S = S_{t-n}, ..., S_t, ..., S_{t+n}$, where $S_i = s_i^1, ..., s_i^{k_i}$ is the sequence of all associated senses in context of $w_i \in W$. Note that S_i might be empty if the word w_i does not have any associated sense. In our model each target word takes as context both its surrounding words and all the senses associated with them. In contrast to the original CBOW architecture, where the training criterion is to correctly classify w_t, our approach aims to predict the word w_t and its set S_t of associated senses. This is equivalent to minimize the following loss function:

$$E = -\log(p(w_t|W^t, S^t)) - \sum_{s \in S_t} \log(p(s|W^t, S^t))$$

where $W^t = w_{t-n}, ..., w_{t-1}, w_{t+1}, ..., w_{t+n}$ and $S^t = S_{t-n}, ..., S_{t-1}, S_{t+1}, ..., S_{t+n}$. Figure 1 shows the organization of the input and the output layers on a sample training instance. In what follows we present a set of variants of the model on the output and the input layers.

3.1 Output layer alternatives

Both words and senses. This is the default case explained above. If a word has one or more associated senses, these senses are also used as target on a separate output layer.

Only words. In this case we exclude senses as target. There is a single output layer with the

---

[1]Our model can also produce a space of words and synsets as output. The whole process does not change: the only difference is that all synonym senses would be merged into a single synset.
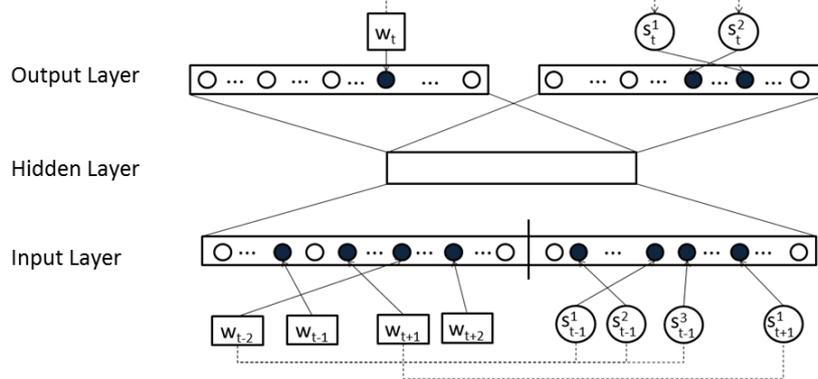
Figure 1: The SW2V architecture on a sample training instance using four context words. Dotted lines represent the link between words and the associated senses in context. In this example, the input layer is composed by a context of two previous words (w_{t-2}, w_{t-1}) and two subsequent words (w_{t+1}, w_{t+2}) with respect to the target word w_t. Two words (w_{t-1}, w_{t+2}) do not have senses associated in context, while w_{t-2}, w_{t+1} have three senses (s_{t-1}^1, s_{t-1}^2, s_{t-1}^3) and one sense associated (s_{t+1}^1) in context, respectively. The output layer is composed by the target word w_t, which has two senses associated (s_t^1, s_t^2) in context.

size of the word vocabulary as in the original CBOW model.

Only senses. In contrast, this alternative excludes words, using only senses as target. In this case, if a word does not have any associated sense, it is not used as a training target instance.

3.2 Input layer alternatives

Both words and senses. Words and their associated senses are included in the input layer and contribute to the hidden state. Both words and senses are updated as a consequence of the backpropagation algorithm.

Only words. In this alternative only the surrounding words contribute to the hidden state. If a word is associated with one or more senses, the update received by the word is also propagated to the weights of the associated senses, even though they are not included in the input layer. This configuration, coupled with the only-words output layer, corresponds exactly to the default CBOW architecture of Word2Vec with an added update step for senses.

Only senses. Words are excluded from the input layer and the target is predicted only from the senses associated with the surrounding words. The weights of the words are updated through the updates of the associated senses

similarly to the previous case. This shared update of words and senses, similarly to the previous alternative, lies on the intuition that a word and a sense are different representations of the same object.

4 Connecting words and senses in context

As mentioned in the previous section, SW2V exploits connections between words and senses in context. One option to obtain such connections could be to take a sense-annotated corpus as input. However, manually annotating large amounts of data is highly expensive and therefore impractical in usual settings. Obtaining sense-annotated data from current off-the-shelf disambiguation and entity linking systems is possible but generally suffers from two major problems. First, supervised systems are hampered by the very same problem of needing large amounts of sense-annotated data. Second, the speed of current disambiguation systems, especially for graph-based unsupervised approaches (Hoffart et al., 2012; Agirre et al., 2014; Moro et al., 2014), could be an obstacle when applied to large corpora.

That is the reason why we propose a simple yet effective unsupervised *shallow word-sense connectivity* algorithm, which can be virtually applied to any given semantic network and it is linear on the corpus size. The main idea of the algorithm is to connect words and senses in context as given by the underlying semantic network. Formally, a

Algorithm 1 Shallow word-sense connectivity

Input: Semantic network (S, E) and text T represented as a bag of words
Output: Set of connected words and senses $T^* \subset T \times S$
1: Set of synsets $S_T \leftarrow \emptyset$
2: **for each** word $w \in T$
3:     $S_T \leftarrow S_T \cup S_w$ (S_w: set of candidate synsets of w)
4: Minimum connections threshold $\theta \leftarrow \frac{|S_T| + |T|}{2\,\delta}$
5: Output set of connections $T^* \leftarrow \emptyset$
6: **for each** $w \in T$
7:     Relative maximum connections $max = 0$
8:     Set of senses associated with w, $C_w \leftarrow \emptyset$
9:     **for each** candidate synset $s \in S_w$
10:        Number of edges $n = |s' \in S_T : (s, s') \in E$ & $\exists w' \in T : w' \neq w$ & $s' \in S_{w'}|$
11:        **if** $n \geq max$ & $n \geq \theta$ **then**
12:            **if** $n > max$ **then**
13:                $C_w \leftarrow \{(w, s)\}$
14:                $max \leftarrow n$
15:            **else**
16:                $C_w \leftarrow C_w \cup \{(w, s)\}$
17:     $T^* \leftarrow T^* \cup C_w$
18: **return** Output set of connected words and senses T^*

---

corpus and a semantic network are taken as input and a set of connected words and senses is produced as output. We define a semantic network as a graph (S, E) where the set S is composed by synsets (nodes) and E is composed by a set of semantically connected synset pairs (edges).

Shallow word-sense connectivity algorithm. Algorithm 1 describes how to connect words and senses in a given text (sentence or paragraph) T. First, we gather in a set S_T all candidate synsets of the words (including multiword expressions up to trigrams) in T (lines 1 to 3). Second, for each candidate synset s we calculate the number of synsets which are connected with s in the semantic network and are included in S_T, excluding connections of synsets which only appear as candidates of the same word (lines 5 to 10). Finally, each word is associated with its top candidate synset(s) according to their number of connections in context, provided that its/their number of connections exceeds a threshold $\theta = \frac{|S_T| + |T|}{2\,\delta}$ (lines 11 to 17). This parameter aims to retain relevant connectivity across senses, as only senses above the threshold will be connected to words in the output corpus. θ is directly proportional to a parameter δ, and to the average text length and number of candidate synsets within the text[2] (line 4 in the Algorithm).

The complexity of the proposed algorithm is

$N + (N \times \alpha)$, where N is the number of words of the training corpus and α is the average polysemy degree of a word in the corpus according to the input semantic network. Considering that non-content words are not taken into account (i.e. polysemy degree 0) and that the average polysemy degree of words in current lexical resources (e.g. WordNet or BabelNet) does not exceed a small constant (3) in any language, we could safely assume that the algorithm is linear in the size of the training corpus. Hence, the training of senses from raw text, considering current predictive models, is maintained with respect to training words only. This enables a fast training on large amounts of text corpora, as opposed to current unsupervised disambiguation algorithms. Additionally, as we will show in Section 5.2, this algorithm does not only speed up significantly the training phase, but also leads to more accurate results.

Note that in our algorithm a word is allowed to have more than one sense associated. In fact, current lexical resources like WordNet (Miller, 1995) or BabelNet (Navigli and Ponzetto, 2012) are hampered by the high granularity of their sense inventories (Hovy et al., 2013). In Section 6.2 we show how our sense embeddings are particularly suited to deal with this issue.

5 Analysis of Model Components

In this section we analyze the different components of SW2V, including the nine model configurations and the connection between words and senses in context. In what follows we describe the common analysis setting:

- **Training model and hyperparameters.** For evaluation purposes, we use the CBOW model of Word2Vec with standard hyperparameters: the dimensionality of the vectors is set to 300 and the window size to 8, and hierarchical softmax is used for normalization. These hyperparameter values are set across all experiments.

- **Corpus and semantic network.** We use a 1.5GB-size corpus from the UMBC project, which is composed by English paragraphs extracted from the web[3]. As semantic network we use BabelNet[4], a large multilin-

---

		Output										
		Words				Senses				Both			
		WS-Sim		RG-65		WS-Sim		RG-65		WS-Sim		RG-65	
		r	ρ	r	ρ	r	ρ	r	ρ	r	ρ	r	ρ
Input	**Words**	0.49	0.48	0.65	0.66	0.56	0.56	0.67	0.67	0.54	0.53	0.66	0.65
	Senses	0.69	0.69	0.70	0.71	0.69	0.70	0.70	**0.74**	**0.72**	**0.71**	**0.71**	**0.74**
	Both	0.60	0.65	0.67	0.70	0.62	0.65	0.66	0.67	0.65	**0.71**	0.68	0.70

Table 1: Pearson (r) and Spearman (ρ) correlation performance of the nine configurations of SW2V

gual semantic network with over 350 million semantic connections, integrating resources such as Wikipedia and WordNet. We chose BabelNet for its wide coverage of named entities and lexicographic knowledge and its immediate scaling up to arbitrary languages.

- **Benchmark.** Word similarity has been one of the most popular benchmarks for *in-vitro* evaluation of vector space models (Pennington et al., 2014; Levy et al., 2015). For the analysis we use two word similarity datasets: the similarity portion (Agirre et al., 2009, WS-Sim) of the WordSim-353 dataset (Finkelstein et al., 2002) and RG-65 (Rubenstein and Goodenough, 1965). In order to compute the similarity of two words using our sense embeddings, we apply the standard closest senses strategy, using cosine similarity as comparison measure between senses.

5.1 Model configurations

In this section we analyze the different configurations of our model with respect to the input and the output layer on a word similarity experiment. Recall from Section 3 that our model could have words, senses or both in either the input and output layers. Table 1 shows the results of all nine configurations on the WS-Sim and RG-65 datasets.

As shown in Table 1, the best configuration according to both Spearman and Pearson correlation measures is the configuration which has only senses in the input layer and both words and senses in the output layer. In fact, taking only senses as input seems to be consistently the best alternative for the input layer. Our intuition is that the knowledge learned from both the co-occurrence information and the the semantic network is more balanced with this input setting. For instance, in the case of including both words and senses in the input layer, the co-occurrence information learned by the network would be duplicated for both words and senses.

5.2 Disambiguation / Shallow word-sense connectivity algorithm

In this section we evaluate the impact of our *shallow word-sense connectivity algorithm* (Section 4) by testing our model directly taking a pre-disambiguated text as input. In this case the network exploits the connections between each word and its disambiguated sense in context. For this comparison we used Babelfy[5] (Moro et al. 2014), a state-of-the-art graph-based disambiguation and entity linking system based on BabelNet. We compare to both the default Babelfy system which uses the *Most Common Sense* (MCS) heuristic as a back-off strategy and, following (Iacobacci et al., 2015), we also include a version in which only instances above the Babelfy default confidence threshold are disambiguated (i.e., the MCS back-off strategy is disabled). We will refer to this latter version as Babelfy* and report the best configuration of each strategy according to our analysis.

Table 2 shows the results of our model using the three different strategies on RG-65 and WS-Sim. Our shallow word-sense connectivity algorithm achieves the best overall results. We believe that these results are due to the semantic connectivity ensured by our algorithm and to the possibility of associating words with more than one sense, which seems beneficial for training, making it more robust to possible disambiguation errors and to the sense granularity issue (Erk et al., 2013). The results are especially significant considering that our algorithm took a tenth of the time needed by Babelfy to disambiguate the corpus.

6 Evaluation

We perform a qualitative and quantitative evaluation of important features of SW2V in three different tasks. First, in order to compare our model against standard word-based approaches, we evaluate our system in the word similarity task (Sec-

---

[5] http://babelfy.org

	WS-Sim		RG-65	
	r	ρ	r	ρ
Shallow	**0.72**	**0.71**	**0.71**	**0.74**
Babelfy	0.65	0.63	0.69	0.70
Babelfy*	0.63	0.61	0.65	0.64

Table 2: Pearson (r) and Spearman (ρ) correlation performance of SW2V integrating our *shallow* word-sense connectivity algorithm (default), Babelfy, or Babelfy*.

tion 6.1). Second, we measure the quality of our sense embeddings in a sense-specific application: sense clustering (Section 6.2). Finally, we evaluate the coherence of our unified vector space by measuring the inter-connectivity of word and sense embeddings (Section 6.3). As future work we plan to develop a WSD and Entity Linking system around this shared space of word and sense embeddings to test our approach in tasks such as contextual word similarity (Huang et al., 2012) and other downstream tasks.

Comparison systems. We compare with the publicly available pre-trained models of four state-of-the-art sense embeddings approaches: Chen et al. (2014)[6] and AutoExtend[7] (Rothe and Schütze, 2015) based on WordNet, and SensEmbed[8] (Iacobacci et al., 2015) and NASARI[9] (Camacho-Collados et al., 2016) based on BabelNet.

Throughout all the experiments we use the same standard hyperparameters mentioned in Section 5 for both the original Word2Vec implementation and our proposed model SW2V. For our model we use the same optimal configuration according to the analysis of the previous section (only senses as input, and both words and senses as output) for all tasks. As training corpus we took the full 3B-words UMBC webbase corpus described in Section 5 and the Wikipedia corpus[10], used by three of the four comparison systems mentioned above. We use both BabelNet (SW2V$_{BN}$) and WordNet (SW2V$_{WN}$) as semantic networks.

6.1 Word Similarity

In this section we evaluate our sense representations on the standard SimLex-999 (Hill et al., 2015) and MEN (Bruni et al., 2014) word similarity datasets[11]. SimLex and MEN are composed by 999 and 3000 word pairs, respectively. Both datasets contain a fine balance of nouns, verbs and adjectives. As explained in Section 5, we use the closest sense strategy for the word similarity measurement of our model and all sense-based comparison systems. As regards the word embedding models, words are directly compared by using cosine similarity. We also include a *retrofitted* version of the original Word2Vec word vectors (Faruqui et al., 2015, Retrofitting[12]) using WordNet (Retrofitting$_{WN}$) and BabelNet (Retrofitting$_{BN}$) as lexical resources.

Table 3 shows the results of SW2V and all comparison models in SimLex and MEN. SW2V consistently outperforms all sense-based comparison systems using the same corpus, and behaves clearly better than the original Word2Vec trained on the same corpus. Retrofitting decreases the performance of the original Word2Vec on the Wikipedia corpus using BabelNet as lexical resource but significantly improves the original word vectors on the UMBC corpus, obtaining comparable results to our approach. However, while our approach provides a shared space of words and senses, Retrofitting still conflates different meanings of a word into the same vector.

Additionally, we noticed that most of the score divergences between our system and the gold standard scores in SimLex were produced on antonym pairs, which are largely represented in this dataset. In contrast to the consistently low gold similarity scores given to antonym pairs, our system varies its similarity scores depending on the specific nature of the pair[13]. Recent works have managed to obtain significant improvements by tweaking usual word embedding approaches into providing low similarity scores for antonym pairs (Pham et al., 2015; Schwartz et al., 2015; Nguyen et al.,

	System	**Corpus**	**SimLex-999**		**MEN**	
			r	ρ	r	ρ
Senses	SW2V$_{BN}$	UMBC	**0.49**	**0.47**	0.75	0.75
	SW2V$_{WN}$	UMBC	0.46	0.45	**0.76**	**0.76**
	AutoExtend	UMBC	0.47	0.45	0.74	0.75
	AutoExtend	Google-News	0.46	0.46	0.68	0.70
	SW2V$_{BN}$	Wikipedia	0.47	0.43	0.71	0.73
	SW2V$_{WN}$	Wikipedia	0.47	0.43	0.71	0.72
	SensEmbed	Wikipedia	0.43	0.39	0.65	0.70
	Chen et al. (2014)	Wikipedia	0.46	0.43	0.62	0.62
Words	Word2Vec	UMBC	0.39	0.39	0.75	0.75
	Retrofitting$_{BN}$	UMBC	0.47	0.46	0.75	**0.76**
	Retrofitting$_{WN}$	UMBC	0.47	0.46	**0.76**	**0.76**
	Word2Vec	Wikipedia	0.39	0.38	0.71	0.72
	Retrofitting$_{BN}$	Wikipedia	0.35	0.32	0.66	0.66
	Retrofitting$_{WN}$	Wikipedia	0.47	0.44	0.73	0.73

Table 3: Pearson (r) and Spearman (ρ) correlation performance on the SimLex-999 and MEN word similarity datasets.

2016), but this is out of the scope of this paper.

6.2 Sense Clustering

Current lexical resources tend to suffer from the high granularity of their sense inventories (Palmer et al., 2007). In fact, a meaningful clustering of their senses may lead to improvements on downstream tasks (Hovy et al., 2013). In this section we evaluate our synset representations on the Wikipedia sense clustering task. For a fair comparison towards the BabelNet-based comparison systems that use Wikipedia corpus for training, in this experiment we report the results of our model trained on the Wikipedia corpus and using BabelNet as lexical resource only. We consider two Wikipedia sense clustering datasets (500-pair and SemEval) created by Dandala et al. (2013) for the evaluation. In these datasets sense clustering is viewed as a binary classification task in which, given a pair of Wikipedia pages, the system should decide whether clustering them into a single instance or not. To this end, we use our synset embeddings and cluster Wikipedia pages[14] together if their similarity exceeds a threshold γ. In order to set the optimal value of γ, we follow Dandala et al. (2013) and use the first 500-pairs sense clustering dataset for tuning. We set the threshold γ to 0.35, which is the value leading to the highest F-Measure among all values from 0 to 1 with a

---
[14]Since Wikipedia is a resource included in BabelNet, our synset representations are expandable to Wikipedia pages.

	Accuracy	**F-Measure**
SW2V	**87.8**	**63.9**
SensEmbed	82.7	40.3
NASARI	87.0	62.5
Multi-SVM	85.5	-
Mono-SVM	83.5	-
Baseline	17.5	29.8

Table 4: Accuracy and F-Measure percentages of different systems on the SemEval Wikipedia sense clustering dataset.

0.05 step size on the 500-pair dataset. Likewise, we set a threshold for NASARI (0.7) and SensEmbed (0.3) comparison systems.

Finally, we evaluate our approach on the SemEval sense clustering test set. This test set is composed by 925 pairs which were obtained from a set of highly ambiguous words gathered from past SemEval tasks. For comparison, we also include the supervised approach of (Dandala et al., 2013) based on a multi-feature Support Vector Machine classifier trained on an automatically-labeled dataset of the English Wikipedia (Mono-SVM) and Wikipedia in four different languages (Multi-SVM). As naive baseline we include the system which would cluster all given pairs.

Table 4 shows the F-Measure and accuracy results on the SemEval sense clustering dataset. SW2V outperforms all comparison systems ac-

	SemEval-07	SemEval-13
SW2V	**39.9**	**54.0**
AutoExtend	17.6	31.0
Baseline	24.8	34.9

Table 5: F-Measure percentage of different MCS strategies on the SemEval-2007 and SemEval-2013 WSD datasets.

cording to both measures, including the sense representations of NASARI and SensEmbed using the same setup and the same underlying lexical resource. This confirms the capability of our system to accurately capture the semantics of word senses on this sense-specific task.

6.3 Word and sense inter-connectivity

This experiment aims at testing the inter-connectivity between word and sense embeddings in the vector space. There have been previous approaches building a shared space of word and sense embeddings (Chen et al., 2014; Rothe and Schütze, 2015) but, to date, little research has focused on testing the semantic coherence of the vector space. To this end, we evaluate our model on a Word Sense Disambiguation (WSD) task, using our shared vector space of words and senses for obtaining a *Most Common Sense* (MCS) baseline. The MCS baseline is generally integrated into the pipeline of state-of-the-art WSD and Entity Linking systems (Zhong and Ng, 2010; Bordes et al., 2012; Moro et al., 2014) as a back-off strategy. Therefore, a system which automatically identifies the MCS of each word from non-annotated text may be quite valuable, especially for large knowledge resources for which obtaining sense-annotated corpora is highly expensive.

The intuition behind this experiment is that a semantically coherent shared space of words and senses should be able to build a strong baseline for the task, as the MCS of a given word should be closer to the word vector than any other sense. Given a word w, we compute the cosine similarity between w and all its candidate senses, selecting the sense leading to the highest similarity score:

$$MCS(w) = \underset{s \in S_w}{\operatorname{argmax}} \, cos(\vec{w}, \vec{s}) \qquad (1)$$

where $cos(\vec{w}, \vec{s})$ refers to the cosine similarity between the embeddings of w and s. In order to assess the reliability of SW2V against previous models using WordNet as sense inventory, we test our

$company_n^2$ (military unit)　　　$school_n^7$ (group of fish)

AutoExtend	SW2V	AutoExtend	SW2V
$company_n^9$	$battalion_n^1$	school	$schools_n^7$
company	battalion	$school_n^4$	$sharks_n^1$
$company_n^8$	$regiment_n^1$	$school_n^6$	sharks
$company_n^6$	$detachment_n^4$	$school_v^1$	$shoals_n^3$
$company_n^7$	$platoon_n^1$	$school_n^3$	$fish_n^1$
$company_v^1$	$brigade_n^1$	elementary	$dolphins_n^1$
firm	regiment	schools	$pods_n^3$
$business_n^1$	$corps_n^1$	$elementary_a^3$	eels
$firm_n^2$	brigade	$school_n^5$	dolphins
$company_n^1$	platoon	$elementary_a^1$	$whales_n^2$

Table 6: Ten closest word and sense embeddings to the senses $company_n^2$ (military unit) and $school_n^7$ (group of fish).

model on the all-words SemEval-2007 (task 17) (Pradhan et al., 2007) and SemEval-2013 (task 12) (Navigli et al., 2013) WSD datasets. Note that our model using BabelNet as semantic network has a far larger coverage than just WordNet and may be additionally used for Wikification (Mihalcea and Csomai, 2007) and Entity Linking tasks. Since the versions of WordNet vary across datasets and comparison systems, we decided to evaluate the systems on the portion of the datasets covered by all comparison systems (less than 10% of instances were removed from each dataset).

Table 5 shows the results of our system and AutoExtend on the SemEval-2007 and SemEval-2013 WSD datasets. SW2V provides the best MCS results in both datasets. In general, AutoExtend does not accurately capture the predominant sense of a word and performs worse than a baseline that selects the intended sense randomly from the set of all possible senses of the target word.

In fact, AutoExtend tends to create clusters which include a word and all its possible senses. As an example, Table 6 shows the closest word and sense[15] embeddings of our model and AutoExtend to the *military* and *fish* senses of, respectively, *company* and *school*. AutoExtend creates clusters with all the senses of *company* and *school* and their related instances, even if they belong to different domains (e.g., $firm_n^2$ or $business_n^1$ clearly concern the *business* sense of *company)*. Instead, SW2V creates a semantic cluster of word and sense embeddings which are semantically close to the respective $company_n^2$ and $school_n^7$ senses.

---

[15]Notation: $word_n^p$ is the n^{th} sense of *word* in WordNet 3.0 with part of speech p.

7 Conclusion and Future Work

In this paper we proposed SW2V[16], a neural model which learns vector representations for words and senses in a joint training phase by exploiting both text corpora and knowledge from semantic networks. Unlike previous sense-based models which use WordNet as sense inventory, our model achieves a semantically coherent vector space of both words and senses as an emerging feature and is scalable to larger semantic networks. In this work we use WordNet and a high-coverage semantic network like BabelNet for our experiments, but any other semantic network could be integrated into our model. Finally, we showed the advantages of using our approach in various tasks with respect to state-of-the-art word- and sense-based models using the same lexical resource, and showed interesting semantic properties of our unified space of word and sense embeddings.

As future work we plan to apply our representations on different Natural Language Applications such as semantic parsing, Word Sense Disambiguation and Entity Linking. Since the whole approach is language-independent and BabelNet is multilingual, we are also planning to apply the same architecture to other languages and focus on multilingual and cross-lingual applications.

References

Eneko Agirre, Enrique Alfonseca, Keith Hall, Jana Kravalova, Marius Paşca, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and WordNet-based approaches. In *Proc. of NAACL*, pages 19–27.

Eneko Agirre, Oier Lopez de Lacalle, and Aitor Soroa. 2014. Random walks for knowledge-based word sense disambiguation. *Computational Linguistics*, 40(1):57–84.

Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. A Neural Probabilistic Language Model. *The Journal of Machine Learning Research*, 3:1137–1155.

Antoine Bordes, Xavier Glorot, Jason Weston, and Yoshua Bengio. 2012. Joint learning of words and meaning representations for open-text semantic parsing. In *International Conference on Artificial Intelligence and Statistics*, pages 127–135.

Antoine Bordes, Sumit Chopra, and Jason Weston. 2014. Question answering with subgraph embeddings. In *Proc. of EMNLP*, pages 615–620.

Elia Bruni, Nam-Khanh Tran, and Marco Baroni. 2014. Multimodal distributional semantics. *J. Artif. Intell. Res.(JAIR)*, 49(1-47).

José Camacho-Collados, Mohammad Taher Pilehvar, and Roberto Navigli. 2016. Nasari: Integrating explicit knowledge and corpus statistics for a multilingual representation of concepts and entities. *Artificial Intelligence*, 240:36–64.

Xinxiong Chen, Zhiyuan Liu, and Maosong Sun. 2014. A unified model for word sense representation and disambiguation. In *Proceedings of EMNLP*, pages 1025–1035, Doha, Qatar.

Bharath Dandala, Chris Hokamp, Rada Mihalcea, and Razvan C. Bunescu. 2013. Sense clustering using Wikipedia. In *Proc. of RANLP*, pages 164–171, Hissar, Bulgaria.

Katrin Erk, Diana McCarthy, and Nicholas Gaylord. 2013. Measuring word meaning in context. *Computational Linguistics*, 39(3):511–554.

Manaal Faruqui, Jesse Dodge, Sujay K. Jauhar, Chris Dyer, Eduard Hovy, and Noah A. Smith. 2015. Retrofitting word vectors to semantic lexicons. In *Proceedings of NAACL*, pages 1606–1615.

Lev Finkelstein, Gabrilovich Evgeniy, Matias Yossi, Rivlin Ehud, Solan Zach, Wolfman Gadi, and Ruppin Eytan. 2002. Placing search in context: The concept revisited. *ACM Transactions on Information Systems*, 20(1):116–131.

Josu Goikoetxea, Aitor Soroa, Eneko Agirre, and Basque Country Donostia. 2015. Random walks and neural network language models on knowledge bases. In *Proceedings of NAACL*, pages 1434–1439.

Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*.

Johannes Hoffart, Stephan Seufert, Dat Ba Nguyen, Martin Theobald, and Gerhard Weikum. 2012. Kore: keyphrase overlap relatedness for entity disambiguation. In *Proceedings of CIKM*, pages 545–554.

Eduard H. Hovy, Roberto Navigli, and Simone Paolo Ponzetto. 2013. Collaboratively built semi-structured content and Artificial Intelligence: The story so far. *Artificial Intelligence*, 194:2–27.

Eric H. Huang, Richard Socher, Christopher D. Manning, and Andrew Y. Ng. 2012. Improving word representations via global context and multiple word prototypes. In *Proc. of ACL*, pages 873–882, Jeju Island, Korea.

Ignacio Iacobacci, Mohammad Taher Pilehvar, and Roberto Navigli. 2015. Sensembed: Learning sense embeddings for word and relational similarity. In *Proceedings of ACL*, pages 95–105, Beijing, China.

---

[16]Source code and pre-trained embedding models available at `http://lcl.uniroma1.it/sw2v`

Sujay Kumar Jauhar, Chris Dyer, and Eduard Hovy. 2015. Ontologically grounded multi-sense representation learning for semantic vector space models. In *Proceedings of NAACL*.

Richard Johansson and Luis Nieto Pina. 2015. Embedding a semantic network in a word space. In *Proc. of NAACL*, pages 1428–1433.

Omer Levy, Yoav Goldberg, and Ido Dagan. 2015. Improving distributional similarity with lessons learned from word embeddings. *TACL*, 3:211–225.

Rada Mihalcea and Andras Csomai. 2007. Wikify! Linking documents to encyclopedic knowledge. In *Proceedings of the Sixteenth ACM Conference on Information and Knowledge management*, pages 233–242, Lisbon, Portugal.

Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781.

George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.

Andrea Moro, Alessandro Raganato, and Roberto Navigli. 2014. Entity Linking meets Word Sense Disambiguation: a Unified Approach. *TACL*, 2:231–244.

Roberto Navigli and Simone Paolo Ponzetto. 2012. BabelNet: The automatic construction, evaluation and application of a wide-coverage multilingual semantic network. *AIJ*, 193:217–250.

Roberto Navigli, David Jurgens, and Daniele Vannella. 2013. SemEval-2013 Task 12: Multilingual Word Sense Disambiguation. In *Proceedings of SemEval 2013*, pages 222–231.

Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2014. Efficient non-parametric estimation of multiple embeddings per word in vector space. In *Proc. of EMNLP*, pages 1059–1069, Doha, Qatar.

Kim Anh Nguyen, Sabine Schulte im Walde, and Ngoc Thang Vu. 2016. Integrating distributional lexical contrast into word embeddings for antonym-synonym distinction. In *Proc. of ACL*, pages 454–459.

Martha Palmer, Hoa Dang, and Christiane Fellbaum. 2007. Making fine-grained and coarse-grained sense distinctions, both manually and automatically. *Natural Language Engineering*, 13(2):137–163.

Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global vectors for word representation. In *Proceedings of EMNLP*, pages 1532–1543.

Nghia The Pham, Angeliki Lazaridou, and Marco Baroni. 2015. A multitask objective to inject lexical contrast into distributional semantics. In *Proceedings of ACL*, pages 21–26.

Sameer Pradhan, Edward Loper, Dmitriy Dligach, and Martha Palmer. 2007. SemEval-2007 task-17: English lexical sample, SRL and all words. In *Proceedings of SemEval*, pages 87–92.

Joseph Reisinger and Raymond J. Mooney. 2010. Multi-prototype vector-space models of word meaning. In *Proc. of ACL*, pages 109–117.

Sascha Rothe and Hinrich Schütze. 2015. AutoExtend: Extending Word Embeddings to Embeddings for Synsets and Lexemes. In *Proceedings of ACL*, pages 1793–1803, Beijing, China, July.

Herbert Rubenstein and John B. Goodenough. 1965. Contextual correlates of synonymy. *Commun. ACM*, 8(10):627–633.

Hinrich Schütze. 1998. Automatic word sense discrimination. *Computational linguistics*, 24(1):97–123.

Roy Schwartz, Roi Reichart, and Ari Rappoport. 2015. Symmetric pattern based word embeddings for improved word similarity prediction. In *Proc. of CoNLL*, pages 258–267.

Fei Tian, Hanjun Dai, Jiang Bian, Bin Gao, Rui Zhang, Enhong Chen, and Tie-Yan Liu. 2014. A probabilistic model for learning multi-prototype word embeddings. In *Proc. of COLING*, pages 151–160.

David Weiss, Chris Alberti, Michael Collins, and Slav Petrov. 2015. Structured training for neural network transition-based parsing. In *Proceedings of ACL*, pages 323–333, Beijing, China.

Mo Yu and Mark Dredze. 2014. Improving lexical embeddings with semantic knowledge. In *ACL (2)*, pages 545–550.

Zhi Zhong and Hwee Tou Ng. 2010. It Makes Sense: A wide-coverage Word Sense Disambiguation system for free text. In *Proc. of ACL System Demonstrations*, pages 78–83.

Will Y Zou, Richard Socher, Daniel M Cer, and Christopher D Manning. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proc. of EMNLP*, pages 1393–1398.