

# GEOMETRY OF POLYSEMY

**Jiaqi Mu, Suma Bhat, Pramod Viswanath**

Department of Electrical and Computer Engineering

University of Illinois at Urbana Champaign

Urbana, IL 61801, USA

{jiaqimu2,spbhat2,pramodv}@illinois.edu

## ABSTRACT

Vector representations of words have heralded a transformational approach to classical problems in NLP; the most popular example is word2vec. However, a single vector does not suffice to model the polysemous nature of many (frequent) words, i.e., words with multiple meanings. In this paper, we propose a three-fold approach for *unsupervised* polysemy modeling: (a) context representations, (b) sense induction and disambiguation and (c) lexeme (as a word and sense pair) representations. A key feature of our work is the finding that a sentence containing a target word is well represented by a low rank *subspace*, instead of a point in a vector space. We then show that the subspaces associated with a particular sense of the target word tend to *intersect* over a line (one-dimensional subspace), which we use to disambiguate senses using a clustering algorithm that harnesses the Grassmannian geometry of the representations. The disambiguation algorithm, which we call  $K$ -Grassmeans, leads to a procedure to label the different senses of the target word in the corpus – yielding lexeme vector representations, all in an unsupervised manner starting from a large (Wikipedia) corpus in English. Apart from several prototypical target (word,sense) examples and a host of empirical studies to intuit and justify the various geometric representations, we validate our algorithms on standard sense induction and disambiguation datasets and present new state-of-the-art results.

## 1 INTRODUCTION

Distributed representations are embeddings of words in a real vector space, achieved via an appropriate function that models the interaction between neighboring words in sentences (e.g.: neural networks (Bengio et al., 2003; Mikolov et al., 2010; Huang et al., 2012), log-bilinear models (Mnih & Hinton, 2007; Mikolov et al., 2013), co-occurrence statistics (Pennington et al., 2014; Levy & Goldberg, 2014)). Such an approach has been strikingly successful in capturing the (syntactic and semantic) similarity between words (and pairs of words), via simple linear algebraic relations between their corresponding vector representations. On the other hand, the *polysemous* nature of words, i.e., the phenomenon of the same surface form representing multiple senses, is a central feature of the creative process embodying natural languages. For example, a large, tall machine used for moving heavy objects and a tall, long-legged, long-necked bird both share the same surface form “crane”. A vast majority of words, especially frequent ones, are polysemous, with each word taking on anywhere from two to a dozen different senses in many natural languages. For instance, WordNet collects 26,896 polysemous English words with an average of 4.77 senses each (Miller, 1995). Naturally, a single vector embedding does not appropriately represent a polysemous word.

There are currently two approaches to address the polysemy issue (a detailed discussion is in Appendix A): (a) sense specific representation learning (Chen et al., 2014; Rothe & Schütze, 2015), usually aided by hand-crafted lexical resources such as WordNet (Miller, 1995); (b) unsupervised sense induction and sense/lexeme representation learning by inferring the senses directly from text (Huang et al., 2012; Neelakantan et al., 2015; Li & Jurafsky, 2015; Arora et al., 2016b).

Since hand-crafted lexical resources sometimes do not reflect the actual meaning of a target word in a given context (Véronis, 2004) and, more importantly, such resources are lacking in many languages (and their creation draws upon intensive expert human resources), we focus on the second approach

in this paper; such an approach is inherently scalable and potentially plausible with the right set of ideas. Indeed, a human expects the contexts to cue in on the particular sense of a specific word, and successful unsupervised sense representation and sense extraction algorithms would represent progress in the broader area of representation of natural language. Such are the goals of this work.

Firth’s hypothesis – a word is characterized by the company it keeps (Firth, 1957) – has motivated the development of single embeddings for words, but also suggests that multiple senses for a target word could be inferred from its contexts (neighboring words within the sentence). This task is naturally broken into three related questions: (a) how to *represent* contexts (neighboring words of the target word); (b) how to *induce* word senses (partition instances of contexts into groups where the target word is used in the same sense within each group) and (c) how to represent lexemes (word and sense pairs) by vectors.

Existing works address these questions by exploring the latent structure of contexts. In an inspired work, Arora et al. (2016b) hypothesize that the global word representation is a linear combination of its sense representations, models the contexts by a finite number of discourse atoms, and recovers the sense representations via sparse coding of all the vectors of the vocabulary (a global fit). Other works perform a local context-specific sense induction: Li & Jurafsky (2015) introduce a sense-based language model to disambiguate word senses and to learn lexeme representations by incorporating the Chinese restaurant process, Reisinger & Mooney (2010) and Huang et al. (2012) label the word senses by clustering the contexts based on the average of the context word embeddings and learn lexeme representations using the labeled corpus. Neelakantan et al. (2015) retain the representation of contexts by the average of the word vectors, but improves the previous approach by jointly learning the lexeme vectors and the cluster centroids in one shot.

**Grassmannian Model:** We depart from the linear latent models in these prior works by presenting a *nonlinear* (Grassmannian) geometric property of contexts. We empirically observe and hypothesize that the context word representations surrounding a target word reside roughly in a low dimensional *subspace*. Under this hypothesis, a specific sense representation for a target word should reside in all the subspaces of the contexts where the word means this sense. Note that these subspaces need not cluster at all: a word such as “launch” in the sense of “beginning or initiating a new endeavor” could be used in a large variety of contexts. Nevertheless, our hypothesis that large semantic units (such as sentences) reside in low dimensional subspaces implies that the subspaces of all contexts where the target word shares the same meaning should *intersect* non-trivially. This further implies that there exists a direction (one dimensional subspace) that is very close to all subspaces and we treat such an intersection vector as the representation of a group of subspaces. Following this intuition, we propose a three-fold approach to deal with the three central questions posed above:

- **Context Representation:** we define the context for a target word to be a set of left  $W$  and right  $W$  non-functional words of the target word ( $W \approx 10$  in our experiments), including the target word itself, and represent it by a low-dimensional subspace spanned by its context word representations;
- **Sense Induction and Disambiguation:** we induce word senses from their contexts by partitioning multiple context instances into groups, where the target word has the same sense within each group. Each group is associated with a representation – the intersection direction – found via  $K$ -Grassmeans, a novel clustering method that harnesses the geometry of subspaces. Finally, we disambiguate word senses for new context instances using the respective group representations;
- **Lexeme Representation:** the lexeme representations can be obtained by running an off-the-shelf word embedding algorithm on a labeled corpus. We label the corpus through hard decisions (involving *erasure* labels) and soft decisions (probabilistic labels), motivated by analogous successful approaches to decoding of turbo and low density parity check codes (LDPC).

**Experiments:** The lexical aspect of our algorithm (i.e., senses can be induced and disambiguated *individually* for each word) as well as the novel geometry (subspace intersection vectors) jointly allow us to capture subtle shades of senses. For instance, in “Can you hear me? You’re on the **air**. One of the great moments of live television, isn’t it?”, our representation is able to capture the occurrence of “air” to mean “live event on camera”. In contrast, with a global fit such as that in

(Arora et al., 2016b) the senses are inherently limited in the number and type of “discourse atoms” that can be captured.

As a quantitative demonstration of the latent geometry captured by our methods, we evaluate the proposed induction algorithm on standard Word Sense Induction (WSI) tasks. Our algorithm outperforms state-of-the-art on two datasets: (a) SemEval-2010 Task 14 (Manandhar et al., 2010) whose word senses are obtained from OntoNotes (Hovy et al., 2006); and (b) a custom-built dataset built by repurposing the polysemous dataset of (Arora et al., 2016b). In terms of lexeme vector embeddings, our representations have evaluations comparable to state-of-the-art on standard tasks – the word similarity task of SCWS (Huang et al., 2012) – and significantly better on a subset of the SCWS dataset which focuses on polysemous target words and the “police lineup” task of (Arora et al., 2016b). A detailed study of the experiments and quantitative results can be found in Section 3.

## 2 OUR APPROACH

We propose a three-fold approach that includes: (a) context representation, (b) word sense induction and disambiguation, and (c) lexeme representation.

### 2.1 CONTEXT REPRESENTATION

Contexts refer to entire sentences or (long enough) consecutive blocks of words in sentences surrounding a target word. Efficient distributed vector representations for sentences and paragraphs are active topics of research in the literature ((Le & Mikolov, 2014; Tai et al., 2015)), with much emphasis on appropriately relating the individual word embeddings with those of the sentences (and paragraphs) they reside in. The scenario of contexts studied here is similar in the sense that they constitute long semantic units similar to sentences, but different in that we are considering semantic units that all have a *common* target word residing inside them. Instead of a straightforward application of existing literature on sentence (and paragraph) vector embeddings to our setting, we deviate and propose a *non-vector* space representation; such a representation is central to the results of this paper and is best motivated by the following simple experiment.

In this paper, we represent a context  $c$  (as a multiset of words) by a point in the Grassmannian manifold – a subspace (denoted by  $S(c)$ ) spanned by its top  $N$  principle components (denoted by  $\{u_n(c)\}_{n=1}^N$ ). Such representation is motivated by a subspace hypothesis (its empirical validation is given in Appendix B) that the context word representations surrounding a target word reside roughly in a low dimensional *subspace*. A detailed algorithm chart for context representations is provided in Appendix J, for completeness.

### 2.2 SENSE INDUCTION AND DISAMBIGUATION

We now turn to sense induction, a basic task that explores polysemy: in this task, a set of sentences (each containing a common target word) have to be partitioned such that the target word is used in the same sense in all the sentences within each partition. The number of partitions relates to the number of senses being identified for the target word. The geometry of the subset representations plays a key role in our algorithm and we start with this next.

**Geometry of Polysemy** Consider a target monosemous word  $w$  and a context sentence  $c$  containing this word  $w$ . The empirical experiment from the previous section allows us to represent  $c$  by a  $N$ -dimensional subspace of the vectors of the words in  $c$ . Since  $N$  (3~5 in our experiments) is much smaller than the number of words in  $c$  (21, on average), one suspects that the representation associated with  $c$  wouldn’t change very much if the target word  $w$  were expurgated from it, i.e.,  $S(c) \approx S(c \setminus w)$ . Putting these two observations together, one arrives at the following hypothesis, in the context of *monosemous* target words:

*Intersection Hypothesis:* the target word vector  $v(w)$  should reside in the *intersection* of  $S(c \setminus w)$ , where the intersection is over all its contexts  $c$ .

We propose an algorithmic approach to robustly discover the intersection direction involves finding that direction vector that is “closest” to all subspaces; we propose doing so by solving the following

optimization problem:

$$\hat{u}(w) = \arg \min_{\|u\|=1} \sum_{w \in c} d(u, S(c \setminus w))^2, \quad d(u, S) = \sqrt{\|u\|^2 - \sum_{n=1}^N (u^T u_n)^2}, \quad (1)$$

where  $d(v, S)$  is the shortest  $\ell_2$ -distance between  $u$  and subspace  $S$ , and  $u_1, \dots, u_N$  are  $N$  orthonormal basis vectors for subspace  $S$ . It is worth mentioning that we chose  $d(\cdot, \cdot)$  as an  $\ell_2$ -distance for simplicity and there are alternative metrics. Specifically, such metric gives an closed form solution,

$$\hat{u}(w) = \arg \max_{\|u\|=1} \sum_{w \in c} \sum_{n=1}^N (u^T u_n(c \setminus w))^2, \quad (2)$$

which can be solved by taking the first principle component of  $\{u_n(c \setminus w)\}_{w \in c, n=1, \dots, N}$ .

The property that context subspaces of a monosemous word intersect at one direction naturally generalizes to polysemy:

*Polysemy Intersection Hypothesis:* the context subspaces of a polysemous word intersect at different directions for different senses.

A detailed study on validating the hypotheses and the corresponding visualization can be found in Appendix C and Appendix D.

**Sense Induction** We can use the representation of senses by the intersection directions of context subspaces for unsupervised sense induction: supposing the target polysemous word that has  $K$  senses (known ahead of time for now), the goal is to *partition* the contexts associated with this target word into  $K$  groups within each of which the target polysemous word shares the same sense. The fact that two groups of context subspaces, corresponding to different senses, intersect at different directions motivates our geometric algorithm: we note that each one of the contexts belongs to a group associated by the *nearest intersection direction* which serves as a prototype of the group. Part of the task is also to identify the most appropriate intersection direction vectors associated with each group. This task represents a form of unsupervised clustering which can be formalized as the optimization problem below.

Given a target polysemous word  $w$ ,  $M$  contexts  $c_1, \dots, c_M$  containing  $w$  and a number  $K$  indicating the number of senses  $w$  has, we would like to partition the  $M$  contexts into  $K$  sets  $S_1, \dots, S_K$  so as to minimize the distance  $d(\cdot, \cdot)$  of each subspace to the intersection direction of its group,

$$\min_{u_1, \dots, u_K, S_1, \dots, S_K} \sum_{k=1}^K \sum_{c \in S_k} d^2(u_k, S(c \setminus w)). \quad (3)$$

This problem (3) is analogous to the objective of  $K$ -means clustering for vectors and solving it exactly in the worst case can be shown to be NP-hard. We propose a natural algorithm by repurposing traditional  $K$ -means clustering built for vector spaces to the Grassmannian space (a detailed algorithm chart is provided in Appendix K). The difference between  $K$ -Grassmeans and  $K$ -means lies in the maximization step:

- **Initialization:** we randomly initialize  $K$  unit-length vectors  $u_1, \dots, u_K$ .
- **Expectation:** we group contexts based on the distance to each intersection direction:

$$S_k \leftarrow \{c_m : d(u_k, S(c_m \setminus w)) \leq d(u_{k'}, S(c_m \setminus w)) \forall k'\}, \forall k.$$

- **Maximization:** we update the intersection direction for each group with the contexts in the group.

$$u_k \leftarrow \arg \min_u \sum_{c \in S_k} d^2(u, S(c \setminus w)).$$

Note that our algorithm can be run for any one *specific* target word, and makes for efficient *online* sense induction; this is relevant in information retrieval applications where the sense of the query

words may need to be found in real time. To get a qualitative feel for how good  $K$ -Grassmeans is for the sense induction task, we run the following synthetic experiment: we randomly pick  $K$  monosemous words, merge their surface forms to create a single *artificial polysemous* word, collect all the contexts corresponding to the  $K$  monosemous words, replace every occurrence of the  $K$  monosemous words by the single artificial polysemous word. Then we run the  $K$ -Grassmeans algorithm on these contexts with the artificial polysemous word as the target word, so as to recover their original labels (which are known ahead of time, since we merged known monosemous words together to create the artificial polysemous word). Figure 1(a) shows the clustering performances on a realization of the artificial polysemous word made of “monastery” and “phd” (here  $K = 2$ ) and Figure 1(b))shows the clustering performance when  $K = 5$  monosemous words “employers”, “exiled”, “grossed”, “incredible” and “unreleased” are merged together. We repeat the experiment over 1,00 trials with  $K$  varying from 2~8 and the accuracy of sense induction is reported in Figure 1(c). Compared to the baseline algorithm proposed in (Huang et al., 2012; Neelakantan et al., 2015),  $K$ -Grassmean on subspaces outperforms  $K$ -means on average word vectors by 5-6%. We also provide a qualitative study of this algorithm on real data in Appendix E, where we study the semantics of each group for an example target word “columbia”. From these experiments we see that  $K$ -Grassmeans performs very well, qualitatively and quantitatively.




Figure 1: A synthetic experiment to study the performances of  $K$ -Grassmeans: (a) monosemous words: “monastery” and “phd”; (b)  $K = 5$  monosemous words: “employers”, “exiled”, “grossed”, “incredible” and “unreleased”; (c) accuracy versus  $K$ .

A quantitative experiment on large and standardized real datasets (which involves real polysemous target words as opposed to synthetic ones), with a comparison with other algorithms in the literature, is detailed in Section 3, where we see that  $K$ -Grassmeans outperforms state-of-the-art.

**Sense Disambiguation** Having the intersection directions to represent the senses, we are ready to disambiguate a target word sense in a given context using the learned intersection directions specific to this target word: for a new context instance for a polysemous word, the goal is to identify which sense this word means in the context. Our approach is three-fold: represent the context by a low dimensional subspace  $S(c \setminus w)$  approximation of the linear span of the word embeddings of non-functional words in the context, find the orthogonal projection distance between the intersection vector  $u_k(w)$  and the context subspace, and finally output  $k^*$  that minimizes the distance, i.e.,

$$k^* = \arg \min_k d(u_k(w), S(c \setminus w)). \quad (4)$$

We refer to (4) as a *hard* decoding of word senses since this outputs a deterministic label. At times, it makes sense to consider a *soft* decoding algorithm where the output is a probability distribution. The probability that  $w$  takes  $k$ -th sense given the context  $c$  is defined via,

$$P(w, c, k) = \frac{\exp(-d(u_k(w), S(c \setminus w)))}{\sum_{k'} \exp(-d(u_{k'}(w), S(c \setminus w)))}. \quad (5)$$

Here we calculate the probability as a monotonic function of the cosine distance between the intersection vector  $u_k(w)$  and the context subspace  $S(c \setminus w)$ , inspired by similar heuristics in the literature (Huang et al., 2012). A qualitative study of (4) and (5) is again provided in Appendix E, where we apply (4) and (5) on the target word “columbia” and five sentences as its contexts.

### 2.3 LEXEME REPRESENTATION

Induction and disambiguation are important tasks by themselves, but several downstream applications can use a distributed vector representation of the multiple senses associated with a target word. Just as with word representations, we expect the distributed lexeme representations to have semantic meanings – similar lexemes should be represented by similar vectors.

It seems natural to represent a lexeme  $s_k(w)$  of a given word  $w$  by the intersection vector associated with the  $k$ -th sense group of  $w$ , i.e.,  $u_k(w)$ . Such an idea is supported by an observation that the intersection vector is close to the word representation vector for many monosemous words (a detailed study of this observation is provided in Appendix F). Despite this empirical evidence, somewhat surprisingly, lexeme representation using the intersection vectors turns out to be not such a good idea, and the reason is fairly subtle. It turns out that the intersection vectors are concentrated on a relatively small surface area on the sphere (magnitudes are not available in the intersection vectors) – the cosine similarity between two random intersection vectors among 10,000 intersection vectors (five intersection vectors each for 2,000 polysemous words) is 0.889 on average with standard deviation 0.068. This is quite in contrast to analogous statistics for (global) word embeddings from the word2vec algorithm: the cosine similarity between two random word vectors is 0.134 on average with standard deviation 0.072. Indeed, word vector representations are known to be approximately uniformly scattered on the unit sphere (the so-called isotropy property, see (Arora et al., 2016a)). The intersection vectors cluster together far more and are quite far from being isotropic – yet they are still able to distinguish different senses as shown by the empirical studies and qualitative experiments on prototypical examples above (and also on standard datasets, as seen in Section 3).

Due to this geometric mismatch between word vectors and intersection directions, and corresponding mismatch in linear algebraic properties expected of these distributed representations, it is not appropriate to use the intersection direction as the lexeme vector representation. In this light, we propose to learn the lexeme representations by an alternate (and more direct) procedure: first *label* the polysemous words in the corpus using the proposed disambiguation algorithm from Section 2.2 and then run a standard word embedding algorithm (we use word2vec) on this labeled corpus, yielding lexeme embeddings. There are several possibilities regarding labeling and are discussed next.

**Hard Decodings** We label the corpus using the disambiguation algorithm as in (4). A special label “IDK” representing “I don’t know” is introduced to avoid introducing too many errors during the labeling phase since (a) our approach is based on the bag-of-words model and cannot guarantee to label every sense correctly; for example, “arm” in “the boat commander was also not allowed to resume his career in the Greek Navy due to his missing **arm** which was deemed a factor that could possibly raise enquiries regarding the mission which caused the trauma.” will be labeled as “weapon”); and (b) we are not clear how such errors will affect existing word embedding algorithms.

An “IDK” label is introduced via checking the closest distance between the context subspace and the intersection directions, i.e., let  $u_{k^*}(w)$  be the closest intersection vector of  $w$  to context  $c$ , we will label this instance as  $k^*$  if  $d(u_{k^*}(w), S(c \setminus w)) < \theta$  and “IDK” otherwise, where  $\theta$  is a hyperparameter. A detailed algorithm chart for sense disambiguation and corpus labeling is provided in Appendix L. The “IDK” label includes instances of words that means a rare sense, (for example: “crane” as in stretching the neck), or a confusing sense which requires disambiguation of context words (for example: “book” and “ticket” in “book a flight ticket”). The IDK labeling procedure is inspired by analogous scenarios in reliable communication over noisy channels where the log likelihood ratio of (coded) bits is close to zero and in practice are better labeled as “erasures”, than treating them as informative for the overall decoding task (Cidon et al., 2012).

**Soft Decodings** Another way of labeling is via using the absolute scores of  $K$ -Grassmeans disambiguation for each sense of a target work in a specific context, cf. Equation (5). Soft decoding involves generating a random corpus by sampling one sense for every occurrence of a polysemous word according to its probability distribution from (5). Then lexeme representations are obtained via an application of a standard word embedding algorithm (we use word2vec) on this (random) labeled corpus. Since we only consider words that are frequent enough (i.e., whose occurrence is larger than 10,000), each sense of a polysemous word is sampled enough times to allow a robust lexeme representation with high probability.

Soft decoding benefits in two scenarios: (a) when a context has enough information for disambiguation (i.e., the probability distribution (5) concentrates on one), the random sampling will have a high chance making a correct decision. (b) otherwise (i.e., the probability distribution have more than one peak), the random sampling will have a chance of not making a wrong (irreversible) decision.

### 3 EXPERIMENTS

Throughout this paper we have conducted multiple qualitative and empirical experiments to highlight and motivate the various geometric representations. In this section we evaluate our algorithms (on sense disambiguation method and sense representation) empirically on (standardized) datasets from the literature, allowing us to get a quantitative feel for the performance on large datasets, as well as afford a comparison with other algorithms from the literature.

**Preliminaries** All our algorithms are unsupervised and operate on a large corpus obtained from Wikipedia dated 09/15. We use WikiExtractor ([http://medialab.di.unipi.it/wiki/Wikipedia\\_Extractor](http://medialab.di.unipi.it/wiki/Wikipedia_Extractor)) to extract the plain text. We use the skip-gram model from word2vec Mikolov et al. (2013) as the word embedding algorithm where we use the default parameter setting. We set  $c = 10$  as the context window size and set  $N = 3$  as the rank of PCA. We choose  $K = 2$  and  $K = 5$  in our experiment. For the disambiguation algorithm, we set  $\theta = 0.6$ .

**Baselines** Our main comparisons are with algorithms that conduct unsupervised polysemy disambiguation, specifically the sense clustering method of (Huang et al., 2012), the multi-sense skip gram model (MSSG) of (Neelakantan et al., 2015) with different parameters, and the sparse coding method with a global dictionary of (Arora et al., 2016b). We were able to download the word and sense representations for (Huang et al., 2012; Neelakantan et al., 2015) online, and trained the word and sense representations of (Arora et al., 2016b) on the same corpus as that used by our algorithms.

#### 3.1 SENSE INDUCTION AND DISAMBIGUATION

Word sense induction (WSI) tasks conduct the following test: given a set of context instances containing a target word, one is asked to partition the context instances into groups such that within each group the target word shares the same sense. We test  $K$ -Grassmeans on two datasets – a standard one from the test set of SemEval-2010 (Manandhar et al., 2010) and a custom-built Make-Sense-2016. Appendix G gives a detailed description about the two datasets.

We evaluate the performance of the algorithms on this (disambiguation) task according to standard measures in the literature: V-Measure and paired F-Score; these two evaluation metrics also feature in the SemEval-2010 WSI task (Manandhar et al., 2010). V-measure is an entropy-based external cluster evaluation metric. Paired F-score evaluates clustering performance by converting the clustering problem into a binary classification problem – given two instances, do they belong to the same cluster or not? Both metrics operate on a contingency table  $A = \{a_{tk}\}$ , where  $a_{tk}$  is the number of instances that are manually labeled as  $t$  and algorithmically labeled as  $k$ . A detailed description is given in Appendix M for completeness. Both the metrics range from 0 to 1, and perfect clustering gives a score of 1. Empirical statistics show that V-Measure favors those with a larger number of cluster and paired F-score favors those with a smaller number of cluster.

Table 1: Performances (V-measure and paired F-score (x100)) of Word Sense Induction Task.

algorithms	SemEval-2010			Make-Sense-2016		
	V-Measure	F-score	# cluster	V-Measure	F-score	# cluster
MSSG.300D.6K	6.90	48.43	2.45	14.40	57.91	2.35
NP-MSSG.300D.6K	6.50	52.45	2.56	15.50	55.39	3.05
Huang 2012	10.60	38.05	6.63	15.50	47.40	6.15
average-#cluster=2	5.30	53.39	1.91	15.9	59.33	1.92
average-#cluster=5	11.10	40.66	4.41	20.3	47.80	4.50
subspace-#cluster=2	7.10	<b>57.25</b>	1.86	28.80	<b>64.66</b>	1.98
subspace-#cluster=5	<b>14.40</b>	44.17	4.23	<b>34.30</b>	58.25	4.58

Table 2: Performances (V-measure, paired F-score, supervised recalls (SR) (x100)) of Word Sense Induction Task from SemEval-2010.

algorithms	V-Measure	F-score	SR (60%/40%)	SR (80%/20%)	# cluster
UoY	15.7	49.8	62.4	62.0	11.54
Hermit	16.2	26.7	58.3	57.3	10.78
Duluth-WSI	9	41.1	60.5	59.5	4.15
#cluster=2	7.50	<b>55.0</b>	60.2	60.1	1.88
#cluster=5	<b>16.5</b>	42.8	<b>64.5</b>	<b>65.2</b>	4.23

Table 1 shows the detailed results of our experiments where all algorithms use Wikipedia as the training set, and from where we see that  $K$ -Grassmeans strongly outperforms the others. For completeness, we adapt the same setting from SemEval-2010 by using the given training corpus and compare  $K$ -Grassmeans against the top reporting systems (i.e., UoY (Korkontzelos & Manandhar, 2010), Hermit (Jurgens & Stevens, 2010), Duluth-WSI (Pedersen, 2010)) in (Manandhar et al., 2010) using all four metrics in Table 2. From Table 2 we can see  $K$ -Grassmeans also strongly outperforms participating systems.

### 3.2 LEXEME REPRESENTATION

The key requirement of lexeme representations should be that they have the same properties as word embeddings, i.e., similar lexemes (or monosemous words) should be represented by similar vectors. Hence we evaluate our lexeme representations on a standard word similarity task focusing on context-specific scenarios: the Stanford Contextual Word Similarities (SCWS) dataset (Huang et al., 2012). In addition to word similarity task on SCWS, we also evaluate our lexeme representations on the “police lineup” task proposed in (Arora et al., 2016b).

**Word Similarity on SCWS** The task is as follows: given a pair of target words, the algorithm assigns a measure of similarities between this pair of words. The algorithm is evaluated by checking the degree of agreement between the similarity measure given by the algorithm and the similarity measure given by humans in terms of Spearman’s rank correlation coefficient. Although this SCWS dataset is not meant specifically for polysemy, we repurpose it for our tasks since it asks for the similarity between two words in two given sentential contexts (the contexts presumably provide important clues to the human rater on the senses of the target words) and also because this is a large (involving 2,003 word pairs) and standard dataset in the literature with 10 human rated similarity scores, each rated on an integral scale from 1 to 10. We take the average of 10 human scores to represent the human judgment.

We propose two measures between  $w$  and  $w'$  given their respective contexts  $c$  and  $c'$  – one (denoted by HardSim) is based on the hard decoding algorithm and the other one (denoted by SoftSim) is based on the soft one. HardSim and SoftSim are defined via,

$$\text{HardSim}(w, w', c, c') = d(v_k(w), v_{k'}(w')),$$

where  $k$ , and  $k'$  are the senses obtained via (4),  $v_k(w)$  and  $v_{k'}(w)$  are the lexeme representations for the two senses,  $d(v, v')$  is the cosine similarity between two vectors  $v$  and  $v'$ , i.e.,  $(d(v, v') = v^T v' / \|v\| \|v'\|)$ , and

$$\text{SoftSim}(w, w', c, c') = \sum_k \sum_{k'} P(w, c, k) P(w', c', k') d(v_k(w), v_{k'}(w')),$$

where  $P(w, c, k)$  is the probability that  $w$  takes  $k$ -th sense given the context  $c$  defined in (5).

Table 3 shows the detailed results on this task. Here we conclude that in general our lexeme representations have a similar performance as the state-of-the-art on both soft and hard decisions. It is worth mentioning that the vanilla word2vec representation (which simply ignores the provided contexts) also has a comparable performance – this makes some sense since some of the words in SCWS are monosemous (and hence their meanings are context-dependent).

To separate out the effect of the combination of monosemous and polysemous words in the SCWS dataset, we expurgate monosemous words from the corpus creating a smaller version that we denoted



by SCWS-lite. In SCWS-lite, we only consider the pairs of sentences where *both* target words in one pair have the *same surface form* but different contexts (and hence potentially different senses). SCWS-lite now contains 241 pairs of words, which is roughly 12% of the original dataset. The performances of the various algorithms are detailed in Table 3, from where we see that our representations (and corresponding algorithms) outperform the state-of-the-art and showcases the superior representational performance when it comes to context-specific polysemy disambiguation.

Model	SCWS		SCWS-lite	
	SoftSim	HardSim	SoftSim	HardSim
skip-gram	62.21		-	
MSSG.300D.6K	<b>67.89</b>	55.99	21.30	20.06
NP-MSSG.300D.6K	67.72	58.55	20.14	19.26
Huang 2012	65.7	26.1	-	-
#cluster=2 (hard)	61.03	58.40	9.33	4.97
#cluster=5 (hard)	60.82	53.14	<b>27.40</b>	<b>22.28</b>
#cluster=2 (soft)	63.59	<b>63.67</b>	5.07	6.53
#cluster=5 (soft)	62.46	61.23	16.54	17.83

Table 3: Performance (Spearman’s rank correlation x100) on SCWS task..

**Police Lineup Task** This task is proposed in Arora et al. (2016b) to evaluate the efficacy of their sense representations (via the vector representations of the discourse atoms). The testbed contains 200 polysemous words and their 704 senses, where each sense is defined by eight related words. For example, the “tool/weapon” sense of “axes” is represented by “handle”, “harvest”, “cutting”, “split”, “tool”, and “wood”, “battle”, “chop”. The task is the following: given a polysemous word, the algorithm needs to identify the true senses of this word from a group of  $m$  senses (which includes many distractors) by outputting  $k$  senses from the group. The algorithm is evaluated by the corresponding precision and recall scores.

This task offers another opportunity to test our representations with the others in the literature, and also provide insight into some of those representations themselves. One possible algorithm is to simply use that proposed in Arora et al. (2016b) where we replace their sense representations with ours: Let  $s_k(w)$  denote a sense of  $w$ , and let  $L$  denote the set of words representing a sense. We define a similarity score between a sense representation of  $w$  and a sense set from the  $m$  senses as,  $score(s_k(w), L) = \sqrt{\sum_{w' \in L} (v_k(w)^T v_{w'})^2}$ , and group two senses with highest scores with respect to  $s_k(w)$  for each  $k$ , and then output the top  $k$  senses with highest scores. A detailed algorithm is provided in Appendix H, with a discussion of the potential variants.

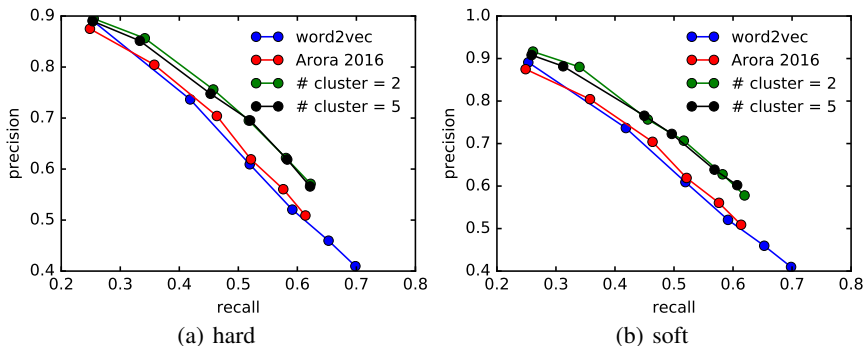


Figure 2: Precision and recall curve in the sense identification task where we let  $m = 20$  and  $k$  vary from 1 to 6.

Figure 2 shows the precision and recall curve in the polysemy test where we let  $m = 20$  and let  $k$  vary from 1 to 6. First, we observe that our representations are uniformly better over the precision-recall curve than the state of the art, although by a relatively small margin. Soft decoding performs slightly better than hard decoding over this task. Second, the surprising finding is that the

baseline we create using vanilla word2vec representations (precise details of this baseline algorithm are provided in Appendix H for completeness) performs as well as the state-of-the-art described in Arora et al. (2016b). A careful look shows that all algorithm outputs (word2vec, ours and those in Arora et al. (2016b)) are highly correlated – they all make correct calls on obvious instances and all make mistakes for confusing instances. We believe that this is because of the following: (a) the algorithm 1 in Arora et al. (2016b) is highly correlated with word2vec since their overall similarity measure uses a linear combination of the similarity measures associated with the sense (discourse atom) vector and the word vector (see Step 6 of Algorithm 1 in Arora et al. (2016b)); (b) word2vec embeddings appear to have an ability to capture two different senses of a polysemous word (as discussed earlier); (c) the instances where the errors occur all seem to be either genuinely subtle or rare in the domain where embeddings were trained (for instance “bat” in the sense of fluttering eyelashes is rare in the Wikipedia corpus, and is one of the error instances).

## 4 CONCLUSION

In this paper, we study the geometry of contexts and polysemy and propose a three-fold approach (entitled  $K$ -Grassmeans) to model target polysemous words in an unsupervised fashion: (a) we represent a context (non-function words surrounding the target word) by a low rank subspace, (b) induce word senses by clustering the subspaces in terms of a distance to an intersection vector and (c) representing lexemes (as a word and sense pair) by labeling the corpus. Our representations are novel and involve nonlinear (Grassmannian) geometry of subspaces and the clustering algorithms are designed to harness this specific geometry. The overall performance of the method is evaluated quantitatively on standardized word sense induction and word similarity tasks and we present new state-of-the-art results. Several new avenues of research in natural language representations arise from the ideas in this work and we discuss a few items in Appendix I.

## REFERENCES

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. A latent variable model approach to pmi-based word embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399, 2016a. ISSN 2307-387X. URL <https://transacl.org/ojs/index.php/tacl/article/view/742>.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. Linear algebraic structure of word senses, with applications to polysemy. *arXiv preprint arXiv:1601.03764*, 2016b.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Jauvin. A neural probabilistic language model. *journal of machine learning research*, 3(Feb):1137–1155, 2003.
- Xinxiong Chen, Zhiyuan Liu, and Maosong Sun. A unified model for word sense representation and disambiguation. In *EMNLP*, pp. 1025–1035. Citeseer, 2014.
- Asaf Cidon, Kanthi Nagaraj, Sachin Katti, and Pramod Viswanath. Flashback: decoupled lightweight wireless control. *ACM SIGCOMM Computer Communication Review*, 42(4):223–234, 2012.
- John Rupert Firth. *Papers in linguistics, 1934-1951*. Oxford University Press, 1957.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. Ontonotes: the 90% solution. In *Proceedings of the human language technology conference of the NAACL, Companion Volume: Short Papers*, pp. 57–60. Association for Computational Linguistics, 2006.
- Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pp. 873–882. Association for Computational Linguistics, 2012.
- David Jurgens and Keith Stevens. Hermit: Flexible clustering for the semeval-2 wsi task. In *Proceedings of the 5th international workshop on semantic evaluation*, pp. 359–362. Association for Computational Linguistics, 2010.

- Ioannis Korkontzelos and Suresh Manandhar. Uoy: Graphs of unambiguous vertices for word sense induction and disambiguation. In *Proceedings of the 5th international workshop on semantic evaluation*, pp. 355–358. Association for Computational Linguistics, 2010.
- Quoc V Le and Tomas Mikolov. Distributed representations of sentences and documents. In *ICML*, volume 14, pp. 1188–1196, 2014.
- Omer Levy and Yoav Goldberg. Neural word embedding as implicit matrix factorization. In *Advances in neural information processing systems*, pp. 2177–2185, 2014.
- Jiwei Li and Dan Jurafsky. Do multi-sense embeddings improve natural language understanding? *arXiv preprint arXiv:1506.01070*, 2015.
- Suresh Manandhar, Ioannis P Klapaftis, Dmitriy Dligach, and Sameer S Pradhan. Semeval-2010 task 14: Word sense induction & disambiguation. In *Proceedings of the 5th international workshop on semantic evaluation*, pp. 63–68. Association for Computational Linguistics, 2010.
- Tomas Mikolov, Martin Karafiát, Lukas Burget, Jan Cernocký, and Sanjeev Khudanpur. Recurrent neural network based language model. In *Interspeech*, volume 2, pp. 3, 2010.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11): 39–41, 1995.
- Andriy Mnih and Geoffrey Hinton. Three new graphical models for statistical language modelling. In *Proceedings of the 24th international conference on Machine learning*, pp. 641–648. ACM, 2007.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. Efficient non-parametric estimation of multiple embeddings per word in vector space. *arXiv preprint arXiv:1504.06654*, 2015.
- Ted Pedersen. Duluth-wsi: Senseclusters applied to the sense induction task of semeval-2. In *Proceedings of the 5th international workshop on semantic evaluation*, pp. 363–366. Association for Computational Linguistics, 2010.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pp. 1532–43, 2014.
- Joseph Reisinger and Raymond J Mooney. Multi-prototype vector-space models of word meaning. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pp. 109–117. Association for Computational Linguistics, 2010.
- Sascha Rothe and Hinrich Schütze. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. *arXiv preprint arXiv:1507.01127*, 2015.
- Kai Sheng Tai, Richard Socher, and Christopher D Manning. Improved semantic representations from tree-structured long short-term memory networks. *arXiv preprint arXiv:1503.00075*, 2015.
- Jean Véronis. Hyperlex: lexical cartography for information retrieval. *Computer Speech & Language*, 18(3):223–252, 2004.

## Appendix: Geometry of Polysemy

### A RELATED WORK

There are two main approaches to model polysemy: one is supervised and uses linguistic resources (Chen et al., 2014; Rothe & Schütze, 2015) and the other is unsupervised inferring senses directly from a large text corpus (Huang et al., 2012; Neelakantan et al., 2015; Li & Jurafsky, 2015; Arora et al., 2016b). Our approach belongs to the latter category.

There are differing approaches to harness hand-crafted lexical resources (such as WordNet): Chen et al. (2014) leverages a “gloss” as a definition for each lexeme, and uses this to model and disambiguate word senses. Rothe & Schütze (2015) models sense and lexeme representations through the *ontology* of WordNet. While the approaches are natural and interesting, they are inherently limited due to (a) the coverage of WordNet: WordNet only covers 26k polysemous words, and the senses for polysemous words are not complete and are domain-agnostic (for example, the mathematical sense for “ring” is missing in WordNet and a majority of occurrences of “ring” mean exactly this sense in the Wikipedia corpus) and (b) the fine-grained nature of WordNet: WordNet senses appear at times too pedantic to be useful in practical downstream applications (for example, “air” in “This show will **air** Saturdays at 2 P.M.” and “air” in “We cannot **air** this X-rated song” are identified to have different meanings).

The unsupervised methods do not suffer from the idiosyncracies of linguistic resources, but are inherently more challenging to pursue since they only rely on the *latent* structures of the word senses embedded inside their contexts. Existing unsupervised approaches can be divided into two categories, based on what aspects of the contexts of target words are used: (a) global structures of contexts and (b) local structures of contexts.

**Global structure:** Arora et al. (2016b) hypothesizes that the global word representation is a linear combination of its sense vectors. This linear algebraic hypothesis is validated by a surprising experiment wherein a single artificial polysemous word is created by merging two random words. The experiment is ingenious and the finding quite surprising but was under a restricted setting: a *single* artificial polysemous word is created by merging only *two* random words. Upon enlargening these parameters (i.e., many artificial polysemous words are created by merging multiple random words) to better suit the landscape of polysemy in natural language, We find the linear-algebraic hypothesis to be fragile: Figure 3 plots the linearity fit as a function of the number of artificial polysemous words created, and also as a function of how many words were merged to create any polysemous word. We see that the linearity fit worsens fairly quickly as the number of polysemous words increases, a scenario that is typical of natural languages.




Figure 3: A synthetic experiment to study the linear algebraic structure of word senses.

The main reason for this effect appears to be that the linearity fit is quite sensitive to the *interaction between the word vectors*, caused by the polysemous nature of the words. The linear algebraic hypothesis is mathematically justified in Section 5 in Arora et al. (2016b) in terms of the RAND-WALK generative model of Arora et al. (2016a) with three extra simplifications. If one were to generalize this proof to handle multiple artificial words at the same time, it appears particularly relevant that the simplification 2 should continue to hold. This simplification step involved the assumption that if  $w_1, \dots, w_n$  be the random words being merged, then (a)  $w_i$  and  $w_j$  do not occur

together in a context window for any  $i \neq j$  and (b) any other word  $\kappa$  can only occur with a single one of the  $w_i$ 's in all context windows. This simplification step clearly no longer holds when  $n$  increases, and especially so when  $n$  nears the size of vocabulary. However, this latter scenario (of  $n$  being the size of the vocabulary) is the very basis of of the sparse-coding algorithm proposed in (Arora et al., 2016b) where the latent structure of the multiple senses is modeled as a corpus of discourse atoms where every atom interacts with all the others.

The experiment, whose results are depicted in Figure 3, is designed to mimic these underlying simplifications of the proof in Arora et al. (2016b): we train word vectors via the skip-gram version of word2vec using the following steps. (a) We initialize the newly generated artificial polysemous words by random vectors; (b) we initialize, and do not update the (two sets of), vector representations of other words  $\kappa$  by the existing word vectors. The embeddings are learnt on the 2016/06/01 Wikipedia dump, tokenized via WikiExtractor ([http://medialab.di.unipi.it/wiki/Wikipedia\\_Extractor](http://medialab.di.unipi.it/wiki/Wikipedia_Extractor)); words that occur less than 1,000 times are ignored; words being merged are chosen randomly in proportion to their frequencies. Due to computational constraints, each instance of mergers is subjected to a single run of the word2vec algorithm.

**Local structure:** Huang et al. (2012); Neelakantan et al. (2015) model a context by the *average* of its constituent word embeddings and use this average vector as a feature to induce word senses by partitioning context instances into groups and to disambiguate word senses for new context instances. Li & Jurafsky (2015) models the senses for a target word in a given context by a Chinese restaurant process, models the contexts also by averaging its constituent word embeddings and then applies standard word embedding algorithms (continuous bag-of-words (CBOW) or skip-gram). Our approach is broadly similar in spirit to these approaches, in that a local lexical-level model is conceived, but we depart in several ways, the most prominent one being the modeling of the contexts as subspaces (and not as vectors, which is what an average of constituent word embeddings would entail).

## B EMPIRICAL VALIDATION OF THE SUBSPACE HYPOTHESIS




Figure 4: An experiment to study the low rank structure of context word embeddings. Histograms of the variance ratio are plotted for rank- $N$  PCAs of the context embeddings.

Given a random word and a set of its contexts (culled from the set of all sentences where the target word appears), we use principle component analysis (PCA) to project the context word embeddings for every context into an  $N$ -dimensional subspace and measure the low dimensional nature of context word embeddings. We randomly sampled 500 words whose occurrence (frequency) is larger than 10,000, extracted their contexts from Wikipedia, and plotted the histogram of the variance ratios being captured by rank- $N$  PCA in Figure 4 for  $N = 3, 4, 5$ . We make the following observations: even rank-3 PCA captures at least 45% of the energy (i.e., variance ratio) of the context word representations and rank-4 PCA can capture at least half of the energy almost surely. As comparison, we note that the average the number of context words is roughly 21 and a rank-4 PCA over a random collection of 21 words would be expected to capture only 20% of the energy (this calculation is justified because word vectors have been observed to possess a spatial isotropy property (Arora et al., 2016a)). All word vectors were trained on the Wikipedia corpus with dimension  $d = 300$  using the skip-gram of word2vec (Mikolov et al., 2013).

This experiment immediately suggests the low-dimensional nature of contexts, and that the contexts be represented in the space of subspaces, i.e., the Grassmannian manifold:

### C EMPIRICAL VALIDATION OF THE INTERSECTION HYPOTHESIS

We illustrate the intersection phenomenon via another experiment. Consider a monosemous word “typhoon” and consider all contexts in Wikipedia corpus where this word appears (there are 14,829 contexts). We represent each context by the rank- $N$  PCA subspace of all the vectors (with  $N = 3$ ) associated with the words in the context and consider their intersection. Each of these subspaces is  $3 \times d$  dimensional (where  $d = 300$  is the dimension of the word vectors). We find that cosine similarity (normalized projection distance) between the vector associated with “typhoon” and each context subspace is very small: the average is 0.693 with standard deviation 0.211. For comparison, we randomly sample 14,829 contexts and find the average is 0.305 with standard deviation 0.041 (a detailed histogram is also provided in Figure 5(a)). This corroborates with the hypothesis that the target word vector is in the intersection of the context subspaces.

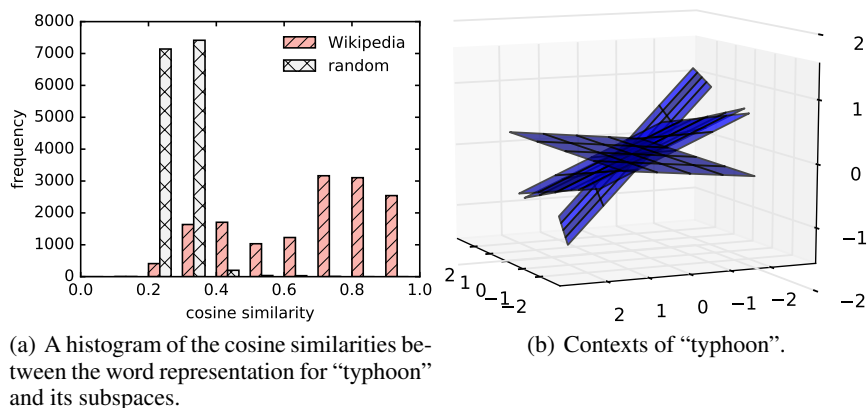


Figure 5: The geometry of contexts for monosemy.

**Visualization of Intersection Hypothesis** A visual representation of this geometric phenomenon is in Figure 5(b), where we have projected the  $d$ -dimensional word representations into 3-dimensional vectors and use these 3-dimensional word vectors to get the subspaces for the four contexts (we set  $N = 2$  for visualization) in Table 4 of the target word “typhoon”, and plot the subspaces as 2-dimensional planes. From Figure 5, we can see that all the context subspaces roughly intersect at a common direction, thus empirically justifying the intersection hypothesis.

“typhoon”
powerful typhoon that affected southern japan in july it was the sixth named storm and second typhoon of the pacific <b>typhoon</b> season originating from an area of low pressure near wake island on july the precursor to maon gradually developed
typhoon ineng was a powerful typhoon that affected southern japan in july it was the sixth named storm and second <b>typhoon</b> of the pacific typhoon season originating from an area of low pressure near wake island on july the precursor
crossing through a gulf of the south china sea patsy weakened to a mph kmh tropical storm before the joint <b>typhoon</b> warning center ceased following the system on the morning of august as it made landfall near the city of
bay were completely wiped out while all airplanes at the local airports were damaged this is the first active pacific <b>typhoon</b> season on record typhoon barbara formed on march and moved west it strengthened briefly to a category three with

Table 4: Contexts containing a monosemous word “typhoon”.

## D EMPIRICAL VALIDATION OF THE POLYSEMY HYPOTHESIS

This intuition behind the polysemy hypothesis is validated by the following experiment, which continues on the same theme as the one done for the monosemous word “typhoon”. Now we study the geometry of contexts for a polysemous word “crane”, which can either mean a large, tall machine used for moving heavy objects or a tall, long-legged, long-necked bird. We list four contexts for each sense of “crane” in Table 5, repeat the experiment as conducted above for the monosemous word “typhoon” and visualize the context subspaces for two senses in Figure 6(a) and 6(b) respectively. Figure 6(c) plots the direction of two intersections. This immediately suggests that the contexts where “crane” stands for a bird intersect at one direction and the contexts where “crane” stands for a machine, intersect at a different direction as visualized in 3 dimensions.




Figure 6: Geometry of contexts for a polysemous word “crane”: (a) all contexts where “crane” means a machine roughly intersect at one direction; (b) all contexts where “crane” means a bird roughly intersect at another direction; (c) two directions representing “crane” as a machine and as a bird.

crane: machine	crane: bird
In June 1979, the anchor towed ashore and lifted by mobile <b>crane</b> into a tank made of concrete built into the ground specifically for the purpose of conserving the anchor.	The sandhill crane (“ <i>Grus canadensis</i> ”) is a species of large <b>crane</b> of North America and extreme northeastern siberia.
The company ran deck and covered lighters, stick lighters, steam cranes and heavy lift <b>crane</b> barges, providing a single agency for Delaware Valley shippers.	Although the grey crowned <b>crane</b> remains common over much of its range, it faces threats to its habitat due to drainage, overgrazing, and
He claimed that his hydraulic <b>crane</b> could unload ships faster and more cheaply than conventional cranes.	The blue crane (“ <i>Anthropoides paradiseus</i> ”), also known as the Stanley <b>crane</b> and the paradise crane, is the national bird of South Africa.
A large pier was built into the harbour to accommodate a heavy lift marine <b>crane</b> which would carry the components into the Northumberland Strait to be installed.	The sarus <b>crane</b> is easily distinguished from other cranes in the region by the overall grey colour and the contrasting red head

Table 5: Contexts containing a polysemous word “crane”.

## E A QUALITATIVE STUDY OF $K$ -GRASSMEANS

To get a qualitative feel of this algorithm on real data, we consider an example target word “columbia” with  $K = 5$  senses. We considered 100K sentences, extracted from the Wikipedia corpus. The goal of sense induction is to partition the set of contexts into 5 groups, so that within each group the target word “columbia” has the same sense. We run  $K$ -Grassmeans for this target word and extract the intersect vectors  $u_1, \dots, u_K$  for  $K = 5$ . One sample sentence for each group is given in Table 6 as an example, from which we can see the first group corresponds to British Columbia in Canada, the second one corresponds to Columbia records, the third one corresponds to

Columbia University in New York, the fourth one corresponds to the District of Columbia, and the fifth one corresponds to Columbia River.

Table 6: Semantics of 5 groups for target word “columbia”.

Group No.	contexts
1	(a) research centres in canada it is located on patricia bay and the former british columbia highway a in sidney british <b>columbia</b> vancouver island just west of victoria international airport the institute is paired with a canadian coast guard base
2	(b) her big break performing on the arthur godfrey show and had since then released a series of successful singles through <b>columbia</b> records hechtlancaaster productions first published the music from their film marty in april and june through cromwell music this
3	(c) fellow at the merrill school of journalism university of maryland college park in she was a visiting scholar at the <b>columbia</b> university center for the study of human rights in haddadcompleted a master of international policy
4	(d) signed into law by president benjamin harrison in march a site for the new national conservatory in the district of <b>columbia</b> was never selected much less built the school continued to function in new york city existing solely from philanthropy
5	(e) in cowlitz county washington the caples community is located west of woodland along caples road on the east shore of <b>columbia</b> river and across the river from columbia city oregon the caples community is part of the woodland school district

The performance of  $K$ -Grassmeans in the context of the target word “columbia” can also be understood via disambiguation. We apply our hard decoding algorithm (4) and our soft decoding algorithm 5 on the five sentences listed in Table 6, the optimal  $k^*$ ’s returned by hard decoding algorithm and the probability distributions  $P(w, c, k)$  returned by the soft decoding algorithm are provided in Table 7. From Table 7 we can see that even though the hard decoding algorithm outputs the correct label, some information is missing if we return a single label  $k^*$ . For example, since we take bag-of-words model in  $K$ -Grassmean, some words (e.g. “school” and “new york city” in context (c) provided in Table 6) suggest that the meaning for “columbia” in this instance might also be Columbia University. The function of those words reflects in the probability distribution returned by the soft decoding algorithm, where we can see the probability that “columbia” in this instance means Columbia University is around 0.13. The misleading result mainly comes from the bag-of-words model, and how to resolve it remains open.

Table 7: Hard decoding and soft decoding for disambiguation of “columbia” in five sentences given in Table 6.

Context No.	hard decoding ( $k^*$ )	soft decoding $P(w, c, k)$				
		$k = 1$	$k = 2$	$k = 3$	$k = 4$	$k = 5$
(a)	1	<b>0.81</b>	0.01	0.01	0.05	0.13
(b)	2	0.02	<b>0.92</b>	0.01	0.04	0.01
(c)	3	0.01	0.00	<b>0.91</b>	0.06	0.01
(d)	4	0.07	0.03	0.13	<b>0.70</b>	0.07
(e)	5	0.04	0.01	0.01	0.05	<b>0.90</b>

## F SIMILARITIES AND DIFFERENCES BETWEEN INTERSECTIONS AND REPRESENTATIONS

We observe another phenomenon that can support our intersection hypothesis – the intersection vector is close to the word representation vector for many monosemous word. We perform an experiment to directly confirm this phenomenon: we randomly sample 500 monosemous words



which occur at least 10,000 times, for each word we compute the intersection vector and check the cosine similarity between the intersection vector and the corresponding word representation of these monosemous words. We find that on average the cosine similarity score is a very high 0.607, with a small standard deviation of 0.095 (a detailed histogram is provided in Figure 7).




Figure 7: A histogram of the cosine similarities between word representations and intersection vectors for monosemous words.

However, this phenomenon does not hold for polysemous words. It turns out, the intersection vectors for polysemous words are concentrated on a relatively small surface area on the sphere. The histogram of the cosine similarity between two random intersection vectors among 10,000 intersections (five intersections for 2,000 polysemous words) and the histogram of the cosine similarity between two random word vectors among 10,000 word embeddings are provided in Figure 8.

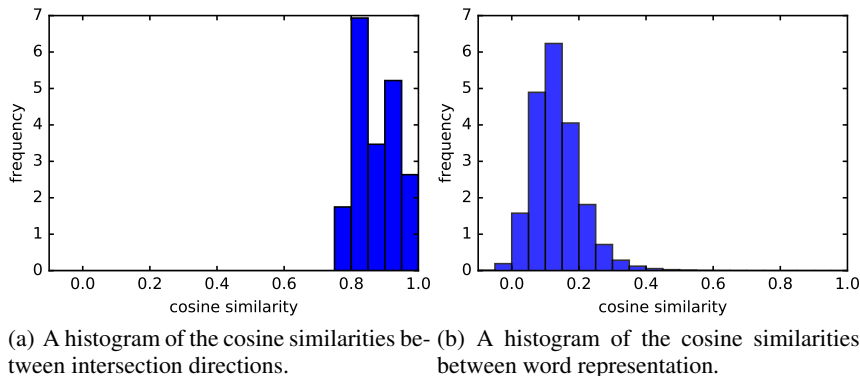


Figure 8: Cosine similarities between intersection directions and word representations respectively.

## G A DETAILED DESCRIPTION OF WSI DATASETS

We test our induction algorithm,  $K$ -Grassmeans, on two datasets – one standard from SemEval-2010 and the other custom-built.

- **SemEval-2010:** The test set of SemEval-2010 shared task 14 (Manandhar et al., 2010) contains 50 polysemous nouns and 50 polysemous verbs whose senses are extracted from OntoNotes (Hovy et al., 2006), and in total 8,915 instances are expected to be disambiguated. The context instances are extracted from various sources including CNN and ABC.
- **Makes-Sense-2016:** Several word senses from SemEval-2010 are too fine-grained in our view (no performance results on tests with native human speakers’ is provided in the literature) – this creates “noise” that reduces the performance of all the algorithms, and the required senses are perhaps not that useful to downstream applications. For example, “guarantee” (as a noun) is labeled to have four different meanings in the following four sentences:

- It has provided a legal **guarantee** to the development of China’s Red Cross cause and connections with the International Red Cross movement, signifying that China’s Red Cross cause has entered a new historical phase.
- Some hotels in the hurricane - stricken Caribbean promise money - back **guarantees**.
- Many agencies roll over their debt , paying off delinquent loans by issuing new loans , or converting defaulted loan **guarantees** into direct loans.
- Litigation consulting isn’t a **guarantee** of a favorable outcome.

However, in general they all mean “a formal promise or pledge”. Towards a more human-interpretable version of the WSI task, we custom-build a dataset whose senses are coarser and clearer. We do this by repurposing a recent dataset created in (Arora et al., 2016b), as part of their “police lineup” task. Our dataset contains 50 polysemous words, together with their senses (on average 2.78 senses per word) borrowed from (Arora et al., 2016b). We generate the testing instances for a target word by extracting all occurrences of it in the Wikipedia Corpus, analyzing its Wikipedia annotations (if any), grouping those which have the same annotations, and finally merging annotations where the target word shares the same sense. Since the senses are quite readily distinguishable from the perspective of native/fluent speakers of English, the disambiguation variability, among the human raters we tested our dataset on, is negligible (this effect is also seen in Figure 6 of (Arora et al., 2016b)).

## H ALGORITHMS FOR POLICE LINEUP TASK

We first introduce the baseline algorithm for word2vec and our algorithm, as given in Algorithm 1 and 2. Both algorithms are motivated by the algorithm in (Arora et al., 2016b).

In our algorithm, the similarity score can be thought as a mean value of the word similarities between a target word  $w$  and a definition word  $w'$  in the given sense  $L$ , we take the power mean with  $p = 2$ . Our algorithm can be adapted to different choice of  $p$ , i.e.,

$$score_p(s_k(w), L) \leftarrow \left( \sum_{w' \in L} |v_k(w)^T v_w|^p \right)^{1/p} - \frac{1}{|V|} \sum_{w'' \in V} \left( \sum_{w' \in L} |v_k(w'')^T v_w|^p \right)^{1/p}$$

Different choice of  $p$  leads to different preferences of the similarities between  $w$  and  $w' \in L$ , generally speaking larger weights put on relevant words with larger  $p$ :

- If we take  $p = 0$ ,  $score_1(w, L)$  turns to be an average of the similarities;
- If we take  $p = \infty$ ,  $score_\infty(w, L)$  measures the similarity between  $w$  and  $L$  by the similarity between  $w$  and the most relevant word  $w' \in L$ , i.e.,

$$score_\infty(s_k(w), L) \leftarrow \max_{w' \in L} |v_k(w)^T v(w')|$$

In our case we take  $p = 2$  to allow enough (but not too much) influence from the most relevant words.

---

**Algorithm 1:** The baseline algorithm with word2vec for Princeton Police Lineup Task.

---

**Input** : a target word  $w$ ,  $m$  senses  $L_1, \dots, L_m$ , word vectors  $v(w')$  for every  $w'$  in the vocabulary  $V$ .

- 1 Compute a similarity score between  $w$  and a sense  $L_i$ ,

$$score(w, L_i) \leftarrow \sqrt{\sum_{w' \in L_i} (v(w)^T v(w'))^2} - \frac{1}{|V|} \sum_{w''} \sqrt{\sum_{w' \in L_i} (v(w'')^T v(w'))^2}$$

- 2 **return** Top  $k$   $L$ 's with highest similarity scores.

**Output:**  $k$  candidate senses.

---

Both our algorithm and the algorithm in (Arora et al., 2016b) do not take into account that one atom represents one sense of the target word, and thus some atoms might generate two senses in the output  $k$  candidates. A more sophisticated algorithm is required to address this issue.

---

**Algorithm 2:** Our algorithm for Princeton Police Lineup Task.

---

**Input** : a target word  $w$ ,  $m$  senses  $L_1, \dots, L_m$ , lexeme representations  $v_k(w)$  for every sense of  $w$ , and word representations  $v(w')$  for  $w'$  in the vocabulary.

```

1 candidates  $\leftarrow$  list()
2 scores  $\leftarrow$  list()
3 for every sense  $s_k(w)$  of  $w$  do
4   for every  $i = 1, 2, \dots, m$  do
5     Compute a similarity score between a lexeme  $s_k(w)$  and a sense  $L_i$ ,
           
$$\text{score}(s_k(w), L) \leftarrow \sqrt{\sum_{w' \in L} (v_k(w)^T v_{w'})^2} + \sqrt{\sum_{w' \in L} (v(w)^T v(w'))^2} - \frac{1}{|V|} \sum_{w'' \in V} \sqrt{\sum_{w' \in L} (v(w'')^T v(w'))^2},$$

6   end
7    $L_{i_1}, L_{i_2} \leftarrow$  top 2 senses.
8   candidates  $\leftarrow$  candidates +  $L_{i_1} + L_{i_2}$ 
9   scores  $\leftarrow$  scores +  $\text{score}(s_k(w), L_{i_1}) + \text{score}(s_k(w), L_{i_2})$ 
10 end
11 return Top  $k$   $L$ 's with highest similarity scores.
Output:  $k$  candidate senses.

```

---

## I FUTURE WORK

Several new avenues of research in natural language representations arise from the ideas in this work and we discuss a few items below.

- Interactions between Polysemous Words:** One of the findings of this work, via the experiments in Appendix A, is that polysemous words interact with each other in the corpus. One natural way to harness these intersections, and hence to sharpen  $K$ -Grassmeans, is to do an iterative labeling procedure. Both hard decoding and soft decoding (discussed in Section 2.3) can benefit from iterations. In hard decoding, the “IDK” labels may be resolved over multiple iterations since (a) the rare senses can become dominant senses once the major senses are already labeled, and (b) a confusing sense can be disambiguated once the polysemous words in its context are disambiguated. In soft decoding, the probability can be expected to concentrate on one sense since each iteration yields yet more precise context word embeddings. This hypothesis is inspired by the success of such procedures inherent in the message passing algorithms for turbo and LDPC codes in reliable wireless communication which share a fair bit of commonality with the setting of polysemy disambiguation (Richardson & Urbanke, 2008). A quantitative tool to measure the disambiguation improvements from iterations and when to stop the iterative process (akin to the EXIT charts for message passing iterative decoders of LDPC codes (Ten Brink, 2001)) is an interesting research direction, as is a detailed algorithm design of the iterative decoding and its implementation; both of these are beyond the scope of this paper.
- Low Dimensional Context Representation:** A surprising finding of this work is that contexts (sentences) that contain a common target word tend to reside in a low dimensional subspace, as justified via empirical observations in Figure 4. Understanding this geometrical phenomenon in the context of a generative model (for instance, RAND-WALK of (Arora et al., 2016a) is not able to explain this) is a basic problem of interest, with several relevant applications (including language modeling Iyer et al. (1994); Chen & Goodman (1996) and semantic parsing of sentences (textual entailment, for example (Dagan et al., 2006; Bar-Haim et al., 2006; Giampiccolo et al., 2007))). Such an understanding could also provide new ideas for the topical subject of representing sentences and paragraphs (Le & Mikolov, 2014; Wieting et al., 2015; Kenter et al., 2016; Kusner et al., 2015) and eventu-

ally combining with document/topic modeling methods such as Latent Dirichlet Allocation Blei et al. (2003).

3. **Combining Linguistic Resources:** Presently the methods of lexeme representation are either exclusive external resource-based or entirely unsupervised. The unsupervised method of  $K$ -Grassmeans reveals robust clusters of senses (and also provides a soft score measuring the robustness (in terms of how frequent the sense is and how sharply/crisply it is used) of the identified sense). On the other hand, WordNet lists a very detailed number of senses, some frequent and robust but many others very fine grained; the lack of any accompanying metric that relates to the frequency and robustness of this sense (which could potentially be domain/corpus specific) really makes this resource hard to make computational use of, at least within the context of polysemy representations. An interesting research direction would be to try to combine  $K$ -Grassmeans and existing linguistic resources to automatically define senses of multiple *granularities*, along with metrics relating to frequency and robustness of the identified senses.

## J CONTEXT REPRESENTATION ALGORITHM

The pseudocode for context representation (c.f. Section 2.1) is provided in Algorithm 3.

---

**Algorithm 3:** The algorithm for context representation.

---

**Input** : a context  $c$ , word embeddings  $v(\cdot)$ , and a PCA rank  $N$ .

- 1 Compute the first  $N$  principle components of samples  $\{v(w'), w' \in c\}$ ,

$$u_1, \dots, u_N \leftarrow \text{PCA}(\{v(w'), w' \in c\}),$$

$$S \leftarrow \left\{ \sum_{n=1}^N \alpha_n u_n, \alpha_n \in R \right\}$$

**Output:**  $N$  orthonormal basis  $u_1, \dots, u_N$  and a subspace  $S$ .

- 2 ,

---

## K SENSE INDUCTION ALGORITHM

The pseudocode for word sense induction (c.f. Section 2.2) is provided in Algorithm 4.

---

**Algorithm 4:** Word sense induction algorithm with given number of senses.

---

**Input** : contexts  $c_1, \dots, c_M$ , and integer  $K$ .

- 1 Initialize unit length vectors  $u_1, \dots, u_K$  randomly, initialize  $L \leftarrow \infty$ ,

- 2 **while**  $L$  does not converge **do**

- 3 **Expectation step:** assign subspaces to the group which has the closest intersection direction,

$$S_k \leftarrow \{c_m : d(u_k, S(c_m \setminus w)) \leq d(u_{k'}, S(c_m \setminus w)) \forall k'\},$$

- 4 **Maximization Step:** update the new intersection direction that minimize the distances to all subspace in this group as (1).

$$u_k \leftarrow \arg \min_v \sum_{c \in S_k} d^2(u, S(c \setminus w)),$$

$$L \leftarrow \sum_{k=1}^K \sum_{c \in S_k} d^2(u_k, S(c \setminus w)),$$

- 5 **end**

**Output:**  $K$  intersection directions  $v_1, \dots, v_K$ .

---

To ensure good performance, we randomize the intersection directions with multiple different seeds and output the best one in terms of the objective function  $L$ ; this step is analogous to the random initialization conducted in kmeans++ in classical clustering literature (Ostrovsky et al., 2006; Arthur & Vassilvitskii, 2007).

## L SENSE DISAMBIGUATION ALGORITHM

The pseudocode for the word sense disambiguation (c.f. Section 2.2) is provided in Algorithm 5.

---

**Algorithm 5:** The algorithm for sense disambiguation.

---

**Input** : a context  $c$ , word embeddings  $v(\cdot)$ , a PCA rank  $N$ , a set of intersection directions  $u_k(w)$ 's, and a threshold  $\theta$ .

- 1 Compute denoised context subspace,
 
$$S \leftarrow S(c \setminus w),$$
- 2 Compute the distance between  $S(c \setminus w)$  and intersections,
 
$$d_k \leftarrow d(u_k, S),$$
- 3 **if** *hard decoding* **then**
- 4     Get the closest cluster,
 
$$k^* \leftarrow \arg \min d_k,$$
- 5     Check the threshold,
- 6     **if**  $d_{k^*} < \theta$  **then**
- 7         **return**  $k^*$ ;
- 8     **end**
- 9     **return** *IDK*;
- 10 **end**
- 11 **if** *soft decoding* **then**
- 12     Compute the probabilities:
 
$$P(w, c, k) \leftarrow \frac{\exp(-10d(u_k(w), S(c \setminus w)))}{\sum_{k'} \exp(-10d(u_{k'}(w), S(c \setminus w)))},$$
- 13 **end**
- 14 **return**  $P(w, c, k)$  for  $k = 1, \dots, K$ ;

**Output:** An label indicating which sense this word means in this context.

---

## M V-MEASURE AND PAIRED F-SCORE

Clustering algorithms partition  $N$  data points into a set of clusters  $\mathcal{K} = \{1, 2, \dots, K\}$ . Given the ground truth, i.e., another partition of data into another set of clusters  $\mathcal{T} = \{1, 2, \dots, T\}$ , the performance of a clustering algorithm is evaluated based on a contingency table  $A = a_{tk}$  representing the clustering algorithm, where  $a_{tk}$  is the number of data whose ground truth label is  $t$  and algorithm label is  $k$ . There are two intrinsic properties of all desirable evaluation measures:

- The measure should be permutation-invariant, i.e., the measure should be the same if we permute the labels in  $\mathcal{K}$  or  $\mathcal{T}$ .
- The measure should encourage intra-cluster similarity and penalize inter-cluster similarity.

V-Measure (Rosenberg & Hirschberg, 2007) and paired F-Score (Artiles et al., 2009) are two standard measures, the definitions of which are given below.

### M.1 V-MEASURE

V-Measure is an entropy-based metric, defined as a harmonic mean of homogeneity and completeness.

- Homogeneity is satisfied if the data points belong to one algorithm cluster fall into a single ground truth cluster, formally defined as,

$$h = \begin{cases} 1 & \text{if } H(\mathcal{T}) = 0 \\ 1 - \frac{H(\mathcal{T}|\mathcal{K})}{H(\mathcal{T})} & \text{otherwise,} \end{cases}$$

where

$$H(\mathcal{T}|\mathcal{K}) = - \sum_{k=1}^K \sum_{t=1}^T a_{tk} N \log \frac{a_{tk}}{\sum_{t=1}^T a_{tk}},$$

$$H(\mathcal{T}) = - \sum_{t=1}^T \frac{\sum_{k=1}^K a_{tk}}{N} \log \frac{\sum_{k=1}^K a_{tk}}{N}.$$

- Completeness is analogous to homogeneity. Formally this is defined as:

$$c = \begin{cases} 1 & \text{if } H(\mathcal{K}) = 0 \\ 1 - \frac{H(\mathcal{K}|\mathcal{T})}{H(\mathcal{K})} & \text{otherwise,} \end{cases}$$

where

$$H(\mathcal{K}|\mathcal{T}) = - \sum_{t=1}^T \sum_{k=1}^K a_{tk} N \log \frac{a_{tk}}{\sum_{k=1}^K a_{tk}},$$

$$H(\mathcal{K}) = - \sum_{k=1}^K \frac{\sum_{t=1}^T a_{tk}}{N} \log \frac{\sum_{t=1}^T a_{tk}}{N}.$$

Given  $h$  and  $c$ , the V-Measure is their harmonic mean, i.e.,

$$V = \frac{2hc}{h+c}.$$

## M.2 PAIRED F-SCORE

Paired F-score evaluates clustering performance by converting the clustering into a binary classification problem – given two instances, do they belong to the same cluster or not?

For each cluster  $k$  identified by the algorithm, we generate  $\binom{\sum_{t=1}^T a_{tk}}{2}$  instance pairs, and for each ground true cluster, we generate  $\binom{\sum_{k=1}^K a_{tk}}{2}$  instances pairs. Let  $F(\mathcal{K})$  be the set of instance pairs from the algorithm clusters and let  $F(\mathcal{T})$  be the set of instances pairs from ground truth clusters. Precision and recall is defined accordingly:

$$P = \frac{|F(\mathcal{K}) \cap F(\mathcal{T})|}{F(\mathcal{K})},$$

$$R = \frac{|F(\mathcal{K}) \cap F(\mathcal{T})|}{F(\mathcal{T})},$$

where  $|F(\mathcal{K})|$ ,  $|F(\mathcal{T})|$  and  $|F(\mathcal{K}) \cap F(\mathcal{T})|$  can be computed using the matrix  $A$  as below:

$$|F(\mathcal{K})| = \sum_{k=1}^K \binom{\sum_{t=1}^T a_{tk}}{2},$$

$$|F(\mathcal{T})| = \sum_{t=1}^T \binom{\sum_{k=1}^K a_{tk}}{2},$$

$$|F(\mathcal{K}) \cap F(\mathcal{T})| = \sum_{t=1}^T \sum_{k=1}^K \binom{a_{tk}}{2}.$$

## REFERENCES

- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. A latent variable model approach to pmi-based word embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399, 2016a. ISSN 2307-387X. URL <https://transacl.org/ojs/index.php/tacl/article/view/742>.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. Linear algebraic structure of word senses, with applications to polysemy. *arXiv preprint arXiv:1601.03764*, 2016b.
- David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1027–1035. Society for Industrial and Applied Mathematics, 2007.
- Javier Artiles, Enrique Amigó, and Julio Gonzalo. The role of named entities in web people search. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing: Volume 2-Volume 2*, pp. 534–542. Association for Computational Linguistics, 2009.
- Roy Bar-Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. The second pascal recognising textual entailment challenge. In *Proceedings of the second PASCAL challenges workshop on recognising textual entailment*, volume 6, pp. 6–4, 2006.
- David M Blei, Andrew Y Ng, and Michael I Jordan. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022, 2003.
- Stanley F Chen and Joshua Goodman. An empirical study of smoothing techniques for language modeling. In *Proceedings of the 34th annual meeting on Association for Computational Linguistics*, pp. 310–318. Association for Computational Linguistics, 1996.
- Xinxiong Chen, Zhiyuan Liu, and Maosong Sun. A unified model for word sense representation and disambiguation. In *EMNLP*, pp. 1025–1035. Citeseer, 2014.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In *Machine learning challenges. evaluating predictive uncertainty, visual object classification, and recognising textual entailment*, pp. 177–190. Springer, 2006.
- Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pp. 1–9. Association for Computational Linguistics, 2007.
- Eduard Hovy, Mitchell Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. Ontonotes: the 90% solution. In *Proceedings of the human language technology conference of the NAACL, Companion Volume: Short Papers*, pp. 57–60. Association for Computational Linguistics, 2006.
- Eric H Huang, Richard Socher, Christopher D Manning, and Andrew Y Ng. Improving word representations via global context and multiple word prototypes. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers-Volume 1*, pp. 873–882. Association for Computational Linguistics, 2012.
- Rukmini Iyer, Mari Ostendorf, and J Robin Rohlicek. Language modeling with sentence-level mixtures. In *Proceedings of the workshop on Human Language Technology*, pp. 82–87. Association for Computational Linguistics, 1994.
- Tom Kenter, Alexey Borisov, and Maarten de Rijke. Siamese cbow: Optimizing word embeddings for sentence representations. *arXiv preprint arXiv:1606.04640*, 2016.
- Matt J Kusner, Yu Sun, Nicholas I Kolkin, and Kilian Q Weinberger. From word embeddings to document distances. In *Proceedings of the 32nd International Conference on Machine Learning (ICML 2015)*, pp. 957–966, 2015.
- Quoc V Le and Tomas Mikolov. Distributed representations of sentences and documents. In *ICML*, volume 14, pp. 1188–1196, 2014.

- Jiwei Li and Dan Jurafsky. Do multi-sense embeddings improve natural language understanding? *arXiv preprint arXiv:1506.01070*, 2015.
- Suresh Manandhar, Ioannis P Klapaftis, Dmitriy Dligach, and Sameer S Pradhan. Semeval-2010 task 14: Word sense induction & disambiguation. In *Proceedings of the 5th international workshop on semantic evaluation*, pp. 63–68. Association for Computational Linguistics, 2010.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. Efficient non-parametric estimation of multiple embeddings per word in vector space. *arXiv preprint arXiv:1504.06654*, 2015.
- Rafail Ostrovsky, Yuval Rabani, Leonard J Schulman, and Chaitanya Swamy. The effectiveness of Lloyd-type methods for the k-means problem. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06)*, pp. 165–176. IEEE, 2006.
- Tom Richardson and Ruediger Urbanke. *Modern coding theory*. Cambridge University Press, 2008.
- Andrew Rosenberg and Julia Hirschberg. V-measure: A conditional entropy-based external cluster evaluation measure. In *EMNLP-CoNLL*, volume 7, pp. 410–420, 2007.
- Sascha Rothe and Hinrich Schütze. Autoextend: Extending word embeddings to embeddings for synsets and lexemes. *arXiv preprint arXiv:1507.01127*, 2015.
- Stephan Ten Brink. Convergence behavior of iteratively decoded parallel concatenated codes. *IEEE transactions on communications*, 49(10):1727–1737, 2001.
- John Wieting, Mohit Bansal, Kevin Gimpel, and Karen Livescu. Towards universal paraphrastic sentence embeddings. *arXiv preprint arXiv:1511.08198*, 2015.