

Modelling Lexical Ambiguity with Density Matrices

Francois Meyer

University of Amsterdam
francoisrmeyer@gmail.com

Martha Lewis

University of Bristol,
ILLC, University of Amsterdam
martha.lewis@bristol.ac.uk

Abstract

Words can have multiple senses. Compositional distributional models of meaning have been argued to deal well with finer shades of meaning variation known as polysemy, but are not so well equipped to handle word senses that are etymologically unrelated, or homonymy. Moving from vectors to density matrices allows us to encode a probability distribution over different senses of a word, and can also be accommodated within a compositional distributional model of meaning. In this paper we present three new neural models for learning density matrices from a corpus, and test their ability to discriminate between word senses on a range of compositional datasets. When paired with a particular composition method, our best model outperforms existing vector-based compositional models as well as strong sentence encoders.

1 Introduction

An integral part of natural language understanding is the ability to handle lexical ambiguity. Words can have multiple meanings, and the precise meaning of a word only becomes clear when we see it in use - the surrounding context disambiguates it. Word sense disambiguation (WSD) is said to be an ‘AI complete’ problem (Navigli, 2009), that is, a problem that is at least as hard as any other problem in AI, and as such has been the subject of extensive research. Standard approaches treat WSD as a classification problem: given a word in context, the task is to classify it into one of a range of possible senses (Lesk, 1986; Schütze, 1998; Navigli, 2009). A more difficult task is to disambiguate every word in a sentence (Chaplot and Salakhutdinov, 2018). A summary of the state of the art is given in (Raganato et al., 2017). More recently, neural approaches (Hadiwinoto et al., 2019; Huang et al., 2019) use contextualised embeddings

as input to WSD systems, together with knowledge from WordNet. Other neural approaches generate multiple sense vectors per word (Neelakantan et al., 2014; Cheng and Kartsaklis, 2015) or vectors representing a context (Melamud et al., 2016).

Disambiguation can be costly. Each word should be disambiguated with respect to the correct senses of the other words in the sentence, meaning that the computational complexity of the task can become problematic (Chaplot and Salakhutdinov, 2018). Within a compositional framework, the idea is for words to disambiguate automatically in the process of composition (Kintsch, 2001; Mitchell and Lapata, 2008; Baroni et al., 2014; Boleda, 2020).

Within purely vector-based models, the amount of ambiguity that a word vector can represent is limited. Baroni et al. (2014) argue that distributional vectors work well for polysemy, but less so for homonymy. Pieledeu et al. (2015) extend the vector based model of meaning to encompass homonymy by using the notion of a *density matrix*. These can be used to encode a probability distribution over possible meanings of a word in a single representation. Density matrices can also be accommodated within a compositional framework, allowing the ambiguity encoded in the matrix to be resolved via composition.

We use density matrices within a compositional distributional framework to model word and sentence meaning. We propose three new models for building density matrices, based on neural word embedding models. We survey several composition methods for density matrices and evaluate how well our density matrices encode ambiguity and to what extent the composition methods achieve disambiguation, on four disambiguation datasets that test disambiguation in a compositional setting. One of our models (multi-sense Word2DM) emerges as the best model overall. When paired with a particular composition method (Phaser), multi-sense

Word2DM outperforms all other models (including existing baselines and high-performing sentence encoders) on most of the disambiguation tasks.

2 Background

Compositional distributional models come in a range of flavours. [Mitchell and Lapata \(2008\)](#) use simple element-wise operations on vectors. More recently, neural models of composition ([Socher et al., 2012](#); [Bowman et al., 2015](#)) and large networks such as BERT ([Devlin et al., 2019](#)) have been extremely successful. A third flavour is the type-logical tensor-based models of composition ([Baroni and Zamparelli, 2010](#); [Coecke et al., 2010](#); [Paperno et al., 2014](#); [Sadrzadeh and Muskens, 2018](#)). The tensor-based model of composition works as follows. We choose a vector space N for nouns, and another S for sentences, and represent relational words as multilinear maps over these spaces. Intransitive verbs are represented as linear maps $N \rightarrow S$, i.e. matrices in $N \otimes S$. Transitive verbs are represented as maps from two copies of N to S , i.e. order 3 tensors or ‘cubes’ of parameters in $N \otimes S \otimes N$. Composition is performed via *tensor contraction* - an extension of matrix multiplication. Matrices and tensors require many parameters. To alleviate this problem, [Grefenstette and Sadrzadeh \(2011a,b\)](#); [Kartsaklis et al. \(2012\)](#) develop ways of building matrices and tensors from word vectors, some of which are described in section 4.1.1.

We use an extension of the tensor-based approach, based on methods given in [Piedeleu et al. \(2015\)](#); [Bankova et al. \(2018\)](#). Nouns and sentences are represented as *density matrices* and relational words (adjectives, verbs, etc.) are represented as *completely positive maps*, which take density matrices to density matrices.

Representing words with density matrices A density matrix over \mathbb{R}^n is a matrix of the form

$$\rho = \sum_i p_i \vec{v}_i \vec{v}_i^\top \quad (1)$$

where $\{p_i\}_i$ are the probabilities assigned to the vectors $\{\vec{v}_i\}_i$. Density matrices over \mathbb{R}^n are:

1. Symmetric: $\rho^\top = \rho$.
2. Positive semidefinite: $\forall \vec{x} \in \mathbb{R}^n, \vec{x}^\top \rho \vec{x} \geq 0$
3. Unit trace: $\text{tr}(\rho) = 1$

To represent words, we view each word as a probability distribution over senses, and we view the

vectors \vec{v}_i in equation (1) as representing its different senses. For example, the word *bright* could mean *shiny* or *clever*. Suppose that when *bright* is used, it is twice as likely to mean *shiny* as it is to mean *clever*. The density matrix for *bright*, denoted $\llbracket \text{bright} \rrbracket$, is computed as follows:

$$\llbracket \text{bright} \rrbracket = \frac{2}{3} \vec{\text{shiny}} \vec{\text{shiny}}^\top + \frac{1}{3} \vec{\text{clever}} \vec{\text{clever}}^\top$$

Composition with density matrices Since we are working with density matrices, nouns are now maps $N \rightarrow N$, i.e. matrices in $N \otimes N$. Sentences are matrices in $S \otimes S$. This means that intransitive verbs are order-4 tensors that take a noun density matrix as input and give back a sentence density matrix. They live in a space $N \otimes N \otimes S \otimes S$. Transitive verbs are order-6 tensors. Clearly, these spaces get very big very quickly.

To deal with this increase in dimensionality, tricks to create completely positive maps out of density matrices have been proposed ([Lewis, 2019b](#); [Coecke and Meichanetzidis, 2020](#)). This allows composition mechanisms to be specified at the level of density matrices, rather than having to work in the high-order spaces described above. We describe these composition mechanisms in section 3.2.

Other applications of density matrices in NLP include modelling entailment in a compositional setting ([Balkir et al., 2015](#); [Bankova et al., 2018](#); [Sadrzadeh et al., 2018](#); [Lewis, 2019a](#); [Bradley and Vlassopoulos, 2020](#)). [Blacoe et al. \(2013\)](#) also use density matrices to model ambiguity, but in a different setting.

WSD in compositional distributional semantics [Baroni et al. \(2014\)](#) argue that compositional distributional semantic models are particularly able to pick out the more subtle shades of meaning termed *polysemy*. This idea is used in ([Mitchell and Lapata, 2008](#); [Grefenstette and Sadrzadeh, 2011a,b](#); [Kartsaklis et al., 2013](#)), where a range of semantic composition models are tested on datasets built to distinguish different senses of words in context. Neural and distributional models for disambiguation are compared in [Milajevs et al. \(2014\)](#), and the role of ellipsis in disambiguation is investigated in [Wijnholds and Sadrzadeh \(2019\)](#).

3 Methods

3.1 Density Matrix Models

We now introduce the three novel methods that we propose for building density matrices.

3.1.1 BERT2DM

Algorithm 1: BERT2DM training

for each sentence s in a corpus **do**
 | Process s with BERT.
 | Extract and store the contextualised embeddings produced by BERT for the words in s .
end
 Discard all contextualised embeddings corresponding to stop words.
 Apply PCA/SVD to the remaining contextualised embeddings.
for each word v in the vocabulary **do**
 | Compute the density matrix of v as

$$[[v]] = \sum_{i \in \text{ind}(v)} \vec{v}_i \vec{v}_i^\top,$$

where $\text{ind}(v)$ are the indices at which the word v occurs in the corpus and \vec{v}_i is the reduced embedding for v .
end

BERT (Devlin et al., 2019) produces *contextualised embeddings* for words and sentences. Given a sentence, it produces vectors for each word that are specific to that particular context (BERT actually models subword units, but we average the subword embeddings of a word to obtain a contextualised word embedding). BERT2DM uses the contextualised embeddings of BERT to build density matrices that encode multiple senses of a word. BERT is applied to a corpus and the contextualised embeddings for a word w are combined to compute w 's density matrix according to equation (1). The procedure is outlined in algorithm 1.

Since the vectors produced by BERT are fairly large, we apply a dimensionality reduction step (either PCA or SVD) over all content word embeddings before combining to form a density matrix.

We also experiment with clustering the contextual embeddings of a word and applying dimensionality reduction to the cluster centroids instead of the contextualised embeddings. The motivation for this is that clustering contextualised embeddings can produce clusters that correspond to distinct senses (as shown by Wiedemann et al. (2019)).

3.1.2 Word2DM

Word2DM is an extension of Word2Vec (Mikolov et al., 2013a,b) skip-gram with negative sampling

Algorithm 2: Word2DM training

for each word v in the vocabulary **do**
 | Randomly initialise a $n \times m$ matrix B_v .
end
for each target word w_t in the corpus **do**
 | **for** each context word w_c **do**
 | | Sample K negative samples from some noise distribution.
 | | Maximise equation 2 with respect to B_t, B_c , and B_{w_k} for $k = 1, \dots, K$.
 | **end**
end
for each word v in the vocabulary **do**
 | Compute its density matrix as $A_v = B_v B_v^\top$.
end

(SGNS). SGNS modifies word vectors to become closer to words they do occur with, and further away from words they don't occur with (the negative samples). When extending the SGNS algorithm to produce density matrices, we must ensure that the matrices satisfy the conditions resulting from their definition: symmetry, positive semidefiniteness, and unit trace. The first and last are easy to enforce, but preserving positivity is more challenging. To preserve positivity, we utilise the following property of positive semi-definiteness:

Property 3.1. For any matrix B , the product BB^\top is positive semi-definite.

We enforce positive semi-definiteness by training the weights of an intermediary matrix B and computing our density matrix as $A = BB^\top$. By updating the weights of B and computing A we indirectly train positive semi-definite matrices. We modify the training objective of SGNS to maximise the similarity of the density matrices of co-occurring words. The objective function at each target-context prediction is then:

$$J(\theta) = \log \sigma(\text{tr}(A_t A_c)) + \sum_{k=1}^K \log \sigma(-\text{tr}(A_t A_{w_k})) \quad (2)$$

where A_t and A_c are the the density matrices of the target and context words respectively, A_1, A_2, \dots, A_K are the density matrices of K negative samples, and θ is the set of weights of the intermediary matrices B_t, B_c and B_1, B_2, \dots, B_K .

Word2DM is a straightforward extension of Word2Vec for learning density matrices. However,

it turns out that enforcing positive semi-definiteness by introducing intermediary matrices leads to sub-optimal training updates. This can be shown by examining the gradients of equation (2) with respect to the intermediary matrices (derivation in supplementary material).

3.1.3 Multi-sense Word2DM

Multi-sense Word2DM is a modification of Word2DM designed to overcome the gradient issues of Word2DM and to explicitly model ambiguity. Multi-sense Word2DM achieves this through the following changes to Word2DM:

- The columns of the intermediary $n \times m$ matrix B now represent the m different senses of the word. Each sense of a word has its own n -dimensional embedding. The density matrix of a word is still computed as before, and can be expressed in terms of the sense embeddings as

$$A = BB^\top = \sum_{i=1}^m \vec{b}_i \vec{b}_i^\top \quad (3)$$

where $\vec{b}_1, \dots, \vec{b}_m$ are the columns of B corresponding to different senses.

- Each word is also associated with a single vector v_w , which represents it as a context word.
- The following objective function is maximised:

$$J(\theta) = \log \sigma(b_t^\top c_t) + \sum_{k=1}^K \log \sigma(-b_t^\top v_{w_k}) \quad (4)$$

where c_t is the sum of context vectors for all words surrounding the target word and b_t is the the embedding for the relevant sense of the target word. We select b_t by finding the column of B_t most similar to c_t (measured by either cosine similarity or dot product). The full training procedure is outlined in algorithm 3.

Multi-sense Word2DM explicitly models ambiguity by letting the columns of the intermediary matrix represent the different senses of a word. During training the column closest to the context embedding is selected as the relevant sense embedding and only this column is updated. This enables the model to avoid the gradient issues of Word2DM. The objective function being maximised (equation 4) has the same gradient as Word2Vec and therefore does not lead to suboptimal training updates.

3.2 Composition methods

The composition methods we use are based on methods in (Lewis, 2019a; Coecke and Meichanetzidis, 2020). These reduce the high-dimensional representations needed for relation words to composition of density matrices. The relational word is seen as a map that takes nouns as arguments. The composition methods are as follows, using the example of an adjective modifying a noun:

Add : $[[\text{adj}]] + [[\text{noun}]]$

Mult : $[[\text{adj}]] \odot [[\text{noun}]]$

Tensor : $[[\text{adj}]] \otimes [[\text{adj}]] \times [[\text{noun}]]$, where \otimes denotes the Kronecker product and \times denotes tensor contraction.

Phaser : $[[\text{adj}]]^{1/2} [[\text{noun}]] [[\text{adj}]]^{1/2}$

More complex phrases are combined according to their parse. So, a transitive sentence modified with an adjective is composed as (subj(verb(adj obj))). For example, composing the sentence *Bob likes old cars* would consist of the following steps:

$$[[\text{old cars}]] = f([[old]], [[cars]])$$

$$[[\text{likes old cars}]] = f([[likes]], [[old cars]])$$

$$[[\text{Bob likes old cars}]] = f([[Bob]], [[likes old cars]])$$

where the composer f can be substituted by any of the composition methods listed above.

4 Experimental Setup

We evaluate our models on three tasks - word similarity, disambiguation, and a word level ambiguity analysis. The code for training and evaluating our models has been made available at <https://github.com/francois-meyer/lexical-ambiguity-dms>. In this section we introduce the experimental setup used in all of these tasks.

4.1 Baselines

Throughout the experiments we compare our models to existing word and sentence embedding models. For our word embedding baselines we use embeddings produced by three existing models - Word2Vec, GloVe, and FastText. We use the publicly available¹ embeddings trained by Wijnholds and Sadrzadeh (2019). The embeddings are 300-dimensional and were trained on the combined and lemmatised ukWaC and Wackypedia corpora². In the tasks that involve sentence-level semantics (the disambiguation tasks) we compare our models

¹<https://github.com/gijswijnholds/compdisteval-ellipsis>

²wacky.sslmit.unibo.it

Algorithm 3: MS-Word2DM training

for each word w in the vocabulary **do**
 Randomly initialise a $n \times m$ matrix B_w
 and a n -dimensional vector v_w .
end
for each target word w_t in the corpus **do**
 Sum the context vectors of the words
 surrounding w_t within a window of size
 $2l$ to get a context embedding c_t :

$$c_t = \sum_{i=t-l, i \neq t}^{t+l} v_i$$

 Compute the similarity (with either cosine similarity or dot product) of the columns b_1, \dots, b_m of B_t and c_t and extract the most similar column as b_t (the embedding of the relevant sense).
 Sample K negative samples from some noise distribution. Maximise equation 4 with regards to b_t , c_t , and v_{w_k} for $k = 1, \dots, K$.
end
for each word v in the vocabulary **do**
 Compute its density matrix as $A_v = B_v B_v^T$.
end

to existing compositional distributional semantic (CDS) models and neural sentence encoders.

4.1.1 CDS models

CDS models compute a sentence vector as a function of the distributional vectors of the words in the sentence. We use the pre-trained word embeddings of Wijnholds and Sadrzadeh (2019) and compute sentence embeddings by either summing, element-wise multiplying, or applying tensor-based composition. For the phrase *big house* with vectors $\vec{\text{big}}$ and $\vec{\text{house}}$ the different compositional distributional methods will be computed as follows:

Add : $\vec{\text{big}} + \vec{\text{house}}$
Mult : $\vec{\text{big}} \odot \vec{\text{house}}$
Tensor : $\vec{\text{big}} \times \vec{\text{big}}^T \times \vec{\text{house}}$

4.1.2 Sentence encoders

We compare our models to two well-known neural sentence encoders - InferSent (Conneau et al., 2017) and BERT (Devlin et al., 2019). InferSent embeddings are 4096-dimensional, which is much

larger than the word embeddings used in our CDS baselines. We use two pre-trained InferSent models that are publicly available³, referred to as InferSent1 and InferSent2 in our results. We also compare our models to BERT as a sentence encoder. BERT produces an embedding for the entire sentence by adding a special classification token ([CLS]) to the start of every sequence. When BERT is used in a sentence-level task, the [CLS] embedding can be used as a semantic representation for the entire sentence.

Some of our evaluation data contains phrases that are not fully formed sentences. To ensure a fair comparison, we convert all phrases to fully formed sentences for evaluation of the sentence encoders. We added “the” before noun phrases and converted verbs to their present tense form.

4.1.3 Context2DM

We also compare our models to a baseline density matrix model, which we call Context2DM. It is based on the procedure of Schütze (1998) for building multi-sense embeddings. Context2DM builds the density matrix of a word w as follows:

1. Context embeddings are obtained for all the contexts in which w occurs (computed by summing the pre-trained embeddings of all the words that occur around w in a particular context).
2. These context embeddings are clustered (using hierarchical agglomerative clustering for $k = 2, \dots, 10$ and the variance reduction criterion to select the number of clusters) and the resulting cluster centroids subsequently represent the different senses of w .
3. The density matrix of w is computed as the mixture of its sense embeddings i.e. the sum of the outer products of the cluster centroids, normalised to have unit trace.

For the pre-trained word embeddings required for step 1 of the above procedure, we use 17-dimensional word embeddings, trained with the gensim implementation⁴ of Word2Vec on the combined ukWaC+Wackypedia corpus.

4.2 Training

All our density matrices are 17×17 (so 289 parameters). This is closest in size to the 300-dimensional baseline embeddings. We train our

³<https://github.com/facebookresearch/InferSent>

⁴<https://radimrehurek.com/gensim/models/word2vec>

Data set	Format	High similarity example	Low similarity example	# Pairs	# Annotators
ML2008	<u>SV</u>	value slump value decline	value slump value slouch	120	53
GS2011	<u>SVO</u>	people buy house people purchase house	people buy house people bribe house	200	25
GS2012	<u>ASVAO</u>	local family run small hotel local family operate small hotel	local family run small hotel local family move small hotel	194	50
KS2013 -CoNLL	<u>ASVAO</u>	young woman file long nail young woman smooth long nail	young woman file long nail young woman register long nail	194	43

Table 1: Details of the disambiguation data sets (A: adjective, S: subject, V: verb, O: object).

Word2DM and multi-sense Word2DM models on the ukWaC+Wackypedia corpus, consisting of 2.8 billion words. Hyperparameters are provided in the supplementary material. Training these models on a single GPU (Nvidia GeForce GTX 1080 Ti) with 60GB of memory takes around 20 hours per iteration of the training corpus. We present results for four different multi-sense Word2DM models. Two use cosine similarity to compare sense vectors to context vectors, while the other two use the dot product. We also vary the number of senses modelled (the number of columns in the intermediary matrix) between 5 and 10.

We present results for four different BERT2DM models. Two of these cluster the BERT representations into senses before dimensionality reduction, while the other two do not. One of the advantages of clustering the representations is that it reduces the size of the matrix on which dimensionality reduction is applied, so it becomes computationally feasible to train on a larger corpus. We train the unclustered variants on a 10-million word subcorpus of Wackypedia, and the clustered variants on a 20-million word subcorpus. We also vary the dimensionality reduction algorithm between PCA and SVD, to test whether or not centering the contextual embeddings before dimensionality reduction makes any difference. Training BERT2DM takes only a few hours on a 16-core CPU (Intel Xeon Gold 6130) but requires around 4.5GB of memory per 1 million words that it is trained on.

4.3 Data Sets

We test our models on data sets designed to test disambiguation in a compositional setting. Data sets for this task contain sentence pairs with:

- An ambiguous *target* word used in a disambiguating phrase.
- A *landmark* word that has the same meaning as one of the target word’s senses.

- Human judgements of how similar the meaning of the phrase is when the ambiguous word is replaced by the landmark word.

We use four disambiguation data sets to evaluate our models. Three of the four data sets - GS2011 (Grefenstette and Sadrzadeh, 2011a), GS2012, and KS2013-CoNLL (Kartsaklis et al., 2013) - are publicly available⁵, while ML2008 (Mitchell and Lapata, 2008) was obtained privately from the authors of Wijnholds and Sadrzadeh (2019). We show examples and statistics of the data sets in table 1.

5 Results

We introduce each of the evaluation tasks and present our results. For multi-sense Word2DM and BERT2DM we trained four models each, with different hyperparameter settings (as described in section 4.2 and listed in table 2).

	RG	WS	MC	SL	MEN
Word2Vec	.818	.662	.765	.404	.781
GloVe	.826	.571	.732	.399	.773
FastText	.767	.517	.682	.404	.768
Context2DM	.228	.234	.331	.094	.267
Word2DM	.541	.473	.452	.157	.540
MS-Word2DM					
- cos, 5 senses	.768	.556	.670	.290	.680
- cos, 10 senses	.727	.580	.659	.256	.682
- dot, 5 senses	.662	.578	.568	.247	.663
- dot, 10 senses	.679	.596	.612	.281	.663
BERT2DM					
- PCA	.452	.275	.388	.226	.351
- SVD	.428	.317	.392	.234	.327
- PCA + cluster	.383	.219	.381	.153	.251
- SVD + cluster	.315	.205	.294	.091	.207

Table 2: Spearman ρ obtained on word similarity tasks.

⁵<http://compling.eecs.qmul.ac.uk/resources/>

	Verb	Mult	Add	Tensor	Phaser
Word2Vec	.215	.256	.299	.231	
GloVe	.332	.098	.304	.397	
FastText	.181	.281	.198	.137	
BERT			.140		
InferSent1			.207		
InferSent2			.174		
Context2DM	-.069	-.025	-.058	-.025	-.064
Word2DM	-.022	.057	.010	.057	-.007
MS-Word2DM					
- cos, 5 senses	.235	.195	.254	.195	.328
- cos, 10 senses	.248	.204	.210	.204	.217
- dot, 5 senses	.216	.145	.280	.145	.311
- dot, 10 senses	.157	.195	.170	.195	.325
BERT2DM					
- PCA	-.055	-.016	-.101	-.016	-.114
- SVD	.072	.170	.075	.170	.067
- PCA + cluster	.105	.170	.062	.170	.130
- SVD + cluster	.090	-.002	.057	-.002	.049

Table 3: Spearman ρ obtained on ML2008.

5.1 Word Similarity

To validate the quality of our density matrices as general semantic representations we evaluate them on the following standard word similarity data sets: RG (Rubenstein and Goodenough, 1965), WS (Finkelstein et al., 2001), MC (Miller and Charles, 1991), SL (Hill et al., 2015), and MEN (Bruni et al., 2012). We use the evaluation scripts and data sets made publicly available⁶ by Faruqui and Dyer (2014). The results are shown in table 2.

Multi-sense Word2DM performs best out of all the density matrix models, achieving scores comparable to the word embeddings. It substantially improves upon Word2DM, supporting our theoretical findings about Word2DM’s learning issues. Using cosine similarity to select the relevant sense results in slightly better scores. The BERT2DM models perform worst of all our models, but still demonstrate some ability to judge word similarity. There is no clear performance difference between using PCA or SVD for dimensionality reduction. Clustering the BERT representations before dimensionality reduction leads to worse correlation scores.

5.2 Disambiguation

The results we obtain on the disambiguation data sets are presented in tables 3 to 6. In each of these tables our density matrix models are compared to our baselines. Column headings specify the composition methods used to compute the phrase represen-

⁶<https://github.com/mfaruqui/eval-word-vectors>

	Verb	Mult	Add	Tensor	Phaser
Word2Vec	.209	.203	.268	.204	
GloVe	.304	.211	.252	.256	
FastText	.210	.187	.154	.185	
BERT			.266		
InferSent1			.241		
InferSent2			.194		
Context2DM	.037	-.027	.012	-.021	.036
Word2DM	-.059	.001	.019	-.064	-.039
MS-Word2DM					
- cos, 5 senses	.187	.286	.206	.289	.365
- cos, 10 senses	.091	.237	.161	.200	<u>.323</u>
- dot, 5 senses	-.016	.091	-.002	.010	.077
- dot, 10 senses	-.021	.116	.025	.131	.118
BERT2DM					
PCA	-.024	-.107	-.001	-.097	-.046
SVD	-.105	-.304	-.078	-.057	-.074
PCA + cluster	-.056	.030	.008	.007	-.028
SVD + cluster	.045	.013	.031	.037	-.029

Table 4: Spearman ρ obtained on GS2011.

tation. These do not apply to the sentence encoders (BERT and InferSent). The rightmost composition method (Phaser) does not apply to the CDS models. The leftmost column (Verb) compares the semantic representations of the verbs without composition. The best performing models, among the baselines and the density matrices, are indicated in bold. We compare the best-performing density matrix models to the best-performing baseline using a one-sided paired t-test (applying the Bonferroni correction to account for multiple comparisons). We indicate statistically significant improvements over the baseline models, or statistically equivalent scores, by underlining the corresponding scores.

Multi-sense Word2DM is by far the best performing density matrix model. It outperforms all the baseline models on 3 out of the 4 data sets. Among all the composition methods, Phaser most consistently achieves high correlation scores (especially on the more complex data sets). In some cases BERT2DM achieves correlation scores that are comparable to multi-sense Word2DM and the baselines. But in general the BERT2DM density matrices cannot reliably be used to achieve disambiguation.

5.3 Ambiguity Analysis

To investigate to what extent our models encode ambiguity at a word level, we turn to von Neumann entropy (VNE). For a density matrix $\rho = \sum_i p_i \vec{v}_i \vec{v}_i^\top$ the VNE is defined as

$$S(\rho) = -\text{tr}(\rho \ln \rho). \quad (5)$$

	Verb	Mult	Add	Tensor	Phaser
Word2Vec	.270	.155	.334	.260	
GloVe	.413	.219	.297	.231	
FastText	.302	.176	.175	.264	
BERT			.471		
InferSent1			.370		
InferSent2			.372		
Context2DM	-.025	.015	-.063	-.033	.005
Word2DM	-.047	.019	-.092	.043	.074
MS-Word2DM					
- cos, 5 senses	.266	.203	.270	.329	.500
- cos, 10 senses	.214	.208	.263	.304	.397
- dot, 5 senses	-.067	.068	-.103	.022	.126
- dot, 10 senses	-.059	.082	-.040	.068	.126
BERT2DM					
- PCA	.025	.031	.122	.056	.056
- SVD	-.071	.001	-.042	-.037	-.056
- PCA + cluster	-.232	-.073	-.141	-.155	-.232
- SVD + cluster	-.132	-.117	-.154	-.089	-.172

Table 5: Spearman ρ obtained on GS2012.

This can be seen as an extension of Shannon entropy to matrices, and quantifies the amount of information encoded in a density matrix.

We perform two analyses of ambiguity with VNE. First we test whether our density matrices model lexical ambiguity. We do this by investigating whether or not the measured ambiguity of a word’s density matrix correlates with the number of meanings associated with the word. Secondly, we perform a systematic analysis of how ambiguity changes when words are composed into phrases. Using the four disambiguation data sets, we measure the VNE before and after composition, expecting ambiguity to decrease after composition. Something similar was done by [Piedeleu et al. \(2015\)](#), at a smaller scale.

For these experiments, we only report results for one variant of multi-sense Word2DM (cosine similarity, 5 senses) and two variants of BERT2DM (SVD and PCA).

Ambiguity and polysemy To determine the number of senses that a word has, we use WordNet synsets (senses) ([Miller, 1995](#)). We compute the correlation between the VNE of density matrices and the number of synsets associated with words. The correlation coefficients are shown in table 7 and the relationships are plotted in figure 1.

The results show that both BERT2DM and multi-sense Word2DM successfully encode how ambiguous words are. Word2DM exhibits a very low correlation and Context2DM (not plotted) shows none.

	Verb	Mult	Add	Tensor	Phaser
Word2Vec	.201	.222	.194	.190	
GloVe	.152	.154	.127	.083	
FastText	.081	.285	.073	.196	
BERT			.314		
InferSent1			.187		
InferSent2			.190		
Context2DM	-.017	-.074	-.037	-.064	-.006
Word2DM	.149	-.118	.114	-.014	.081
MS-Word2DM					
- cos, 5 senses	.135	.075	.190	.147	<u>.309</u>
- cos, 10 senses	.171	.008	.207	.118	.288
- dot, 5 senses	-.026	.034	.039	.117	.275
- dot, 10 senses	.094	.016	.070	.053	.345
BERT2DM					
- PCA	.084	.082	.089	.081	.187
- SVD	.037	.160	.054	.128	.037
- PCA + cluster	.075	.075	.073	.067	-.011
- SVD + cluster	.036	-.058	.020	.052	.017

Table 6: Spearman ρ obtained on KS2013-CoNLL.

Model	Pearson r	Spearman ρ
Context2DM	-.081	-.043
Word2DM	.053	.112
MS-Word2DM	.296	.295
BERT2DM SVD	.367	.418
BERT2DM PCA	.405	.463

Table 7: Correlation coefficient between the VNE of a word’s density matrix and the number of WordNet synsets to which the word belongs.

Ambiguity and composition VNE allows us to measure how ambiguity evolves through composition. We can compare the ambiguity of a word to the ambiguity of a phrase containing the word. Seeing the context in which a word occurs can reveal which sense of the word is being employed, and should therefore reduce the amount of ambiguity present. The multi-sense Word2DM density matrix for the ambiguous word *run* has a VNE of 1.491. After composition in the sentence *The family run the hotel*, the sentence has a VNE of 0.144, so ambiguity has decreased.

We test whether this phenomenon holds true for our models on the disambiguation data sets, which consist of ambiguous verbs and disambiguating phrases. For each of the data sets we compute the average VNE of the verb density matrices. We compare this to the average VNE of the disambiguating phrases, where the density matrices are composed using different composition methods. The results are presented in tables 8 to 11.

In each of the tables, the leftmost column (Verb) contains the average von Neumann entropy of the

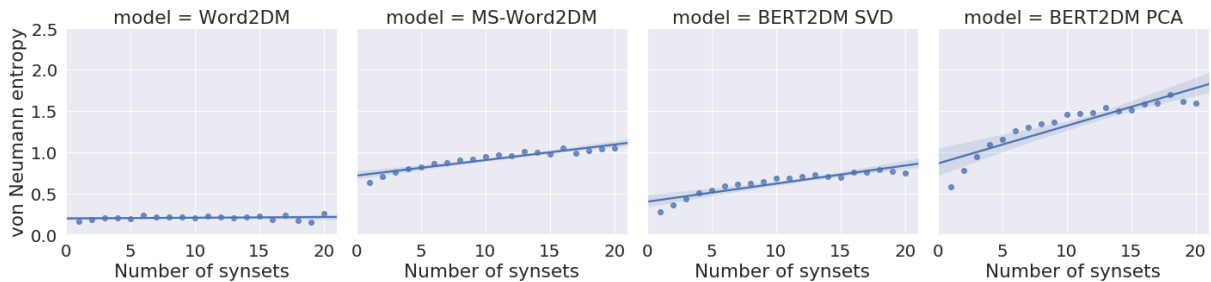


Figure 1: The average VNE of density matrices plotted against the number of WordNet synsets of words.

	Verb	Mult	Add	Tensor	Phaser
Context2DM	0.344	0.354	0.429	0.354	0.007
Word2DM	0.103	0.257	0.707	0.257	0.019
MS-Word2DM	0.953	1.751	1.537	1.751	0.468
BERT2DM*	0.667	0.132	0.938	0.132	0.060

Table 8: Average VNE before and after composition on ML2008. *SVD variant

	Verb	Mult	Add	Tensor	Phaser
Context2DM	0.174	0.187	0.359	0.402	0.000
Word2DM	0.245	0.522	1.096	0.797	0.038
MS-Word2DM	1.444	1.821	1.906	4.012	0.224
BERT2DM*	0.763	0.008	0.961	0.174	0.003

Table 9: Average VNE before and after composition on GS2011. *SVD variant

ambiguous verbs alone, while the other columns contain the average von Neumann entropy after composing the disambiguating phrases with the different composition methods. Cases where the average von Neumann entropy decreases after composition (as expected) are highlighted in bold.

The results are quite similar across the data sets. Phaser emerges as the best method for decreasing ambiguity through composition. This nicely supports the results of the disambiguation experiments, in which Phaser also emerged as the best composition method for disambiguating the meaning of ambiguous words through composition. Besides Phaser, none of the other composition methods reliably decrease the measured ambiguity.

6 Conclusion and Future Work

In this paper we addressed the problem of modelling ambiguity in NLP, and how ambiguous words can be disambiguated in context. We investigated density matrices as semantic representations for modelling ambiguity. Our results confirmed the value of density matrices over vector-based approaches. Equipped with a compositional frame-

	Verb	Mult	Add	Tensor	Phaser
Context2DM	0.167	0.137	0.417	0.387	0.000
Word2DM	0.140	0.623	1.445	0.908	0.005
MS-Word2DM	1.340	1.181	2.043	3.534	0.029
BERT2DM*	0.659	0.000	1.080	0.014	0.000

Table 10: Average VNE before and after composition on GS2012. *SVD variant

	Verb	Mult	Add	Tensor	Phaser
Context2DM	0.233	0.183	0.524	0.537	0.000
Word2DM	0.161	0.604	1.405	0.889	0.010
MS-Word2DM	1.224	1.453	2.032	3.704	0.033
BERT2DM*	0.689	0.000	1.124	0.026	0.000

Table 11: Average VNE before and after composition on KS2013-CoNLL. *SVD variant

work, one of our density matrix models (multi-sense Word2DM) outperformed all other models (including existing compositional models and strong neural baselines) on most of the disambiguation tasks. We also performed a mathematical analysis of the ambiguity encoded by our models. This revealed that the density matrices built by two of our models (multi-sense Word2DM and BERT2DM) reflect true word level ambiguity. We have shown the value in designing neural models that learn density matrices from scratch.

Possible directions for future work includes extending our models to larger datasets and longer sentences, and modifying techniques for differing sentences as in the Word in Context dataset (Pilehvar and Camacho-Collados, 2019) and other WSD tasks. We focused on ambiguity here, but it would be possible to do similar experiments focused on other aspects of meaning such as metaphor or entailment, and examining how these interact with composition.

Acknowledgments

Martha Lewis is supported by NWO Veni grant ‘Metaphorical Meanings for Artificial Agents’.

References

- Esma Balkir, Mehrnoosh Sadrzadeh, and Bob Coecke. 2015. Distributional sentence entailment using density matrices. In *Proceedings of the First International Conference on Theoretical Topics in Computer Science*, volume 9541, pages 1–22.
- Dea Bankova, Bob Coecke, Martha Lewis, and Dan Marsden. 2018. Graded hyponymy for compositional distributional semantics. *Journal of Language Modelling*, 6:225–260.
- Marco Baroni, Raffaella Bernardi, Roberto Zamparelli, et al. 2014. Frege in space: A program for compositional distributional semantics. *Linguistic Issues in language technology*, 9(6):5–110.
- Marco Baroni and Roberto Zamparelli. 2010. Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 1183–1193.
- William Blacoe, Elham Kashefi, and Mirella Lapata. 2013. A quantum-theoretic approach to distributional semantics. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 847–857. Association for Computational Linguistics.
- Gemma Boleda. 2020. Distributional semantics and linguistic theory. *Annual Review of Linguistics*.
- Samuel Bowman, Christopher Potts, and Christopher D Manning. 2015. Recursive neural networks can learn logical semantics. In *Proceedings of the 3rd workshop on continuous vector space models and their compositionality*, pages 12–21.
- Tai-Danae Bradley and Yiannis Vlassopoulos. 2020. Language modeling with reduced densities.
- Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 136–145, Jeju Island, Korea. Association for Computational Linguistics.
- Devendra Singh Chaplot and Ruslan Salakhutdinov. 2018. Knowledge-based word sense disambiguation using topic models. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- Jianpeng Cheng and Dimitri Kartsaklis. 2015. Syntax-aware multi-sense word embeddings for deep compositional models of meaning. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1531–1542, Lisbon, Portugal. Association for Computational Linguistics.
- Bob Coecke and Konstantinos Meichanetzidis. 2020. Meaning updating of density matrices. *ArXiv*, 2001.00862.
- Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. 2010. Mathematical foundations for a compositional distributed model of meaning. *Lambek Festschrift, Linguistic Analysis*, 36.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. Supervised learning of universal sentence representations from natural language inference data. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Jacob Devlin, Ming Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Manaal Faruqi and Chris Dyer. 2014. Community evaluation and exchange of word vectors at wordvectors.org. In *Proceedings of ACL: System Demonstrations*.
- Lev Finkelstein, Evgeniy Gabrilovich, Yossi Matias, Ehud Rivlin, Zach Solan, Gadi Wolfman, and Eytan Ruppin. 2001. Placing search in context: The concept revisited. volume 20, pages 406–414.
- Edward Grefenstette and Mehrnoosh Sadrzadeh. 2011a. Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1394–1404. Association for Computational Linguistics.
- Edward Grefenstette and Mehrnoosh Sadrzadeh. 2011b. Experimenting with transitive verbs in a DisCoCat. In *Proceedings of the GEMS 2011 Workshop on GEometrical Models of Natural Language Semantics*, pages 62–66, Edinburgh, UK. Association for Computational Linguistics.
- Christian Hadiwinoto, Hwee Tou Ng, and Wee Chung Gan. 2019. Improved word sense disambiguation using pre-trained contextualized word representations. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5297–5306, Hong Kong, China. Association for Computational Linguistics.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. SimLex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.

- Luyao Huang, Chi Sun, Xipeng Qiu, and Xuanjing Huang. 2019. [Glossbert: Bert for word sense disambiguation with gloss knowledge](#). *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*.
- Dimitri Kartsaklis, Mehrnoosh Sadrzadeh, and Stephen Pulman. 2012. [A unified sentence space for categorical distributional-compositional semantics: Theory and experiments](#). In *Proceedings of COLING 2012: Posters*, pages 549–558, Mumbai, India. The COLING 2012 Organizing Committee.
- Dimitri Kartsaklis, Mehrnoosh Sadrzadeh, and Stephen Pulman. 2013. [Separating disambiguation from composition in distributional semantics](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning*, pages 114–123, Sofia, Bulgaria. Association for Computational Linguistics.
- Walter Kintsch. 2001. Predication. *Cognitive science*, 25(2):173–202.
- Michael Lesk. 1986. Automatic sense disambiguation using machine readable dictionaries: how to tell a pine cone from an ice cream cone. In *Proceedings of the 5th annual international conference on Systems documentation*, pages 24–26.
- Martha Lewis. 2019a. Compositional hyponymy with positive operators. In *Proceedings of Recent Advances in Natural Language Processing*, pages 638–647.
- Martha Lewis. 2019b. [Modelling hyponymy for disocat](#). Presented at Applied Category Theory Conference 2019.
- Oren Melamud, Jacob Goldberger, and Ido Dagan. 2016. [context2vec: Learning generic context embedding with bidirectional LSTM](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 51–61, Berlin, Germany. Association for Computational Linguistics.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013a. [Efficient estimation of word representations in vector space](#). *CoRR*, abs/1301.3781.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013b. [Distributed representations of words and phrases and their compositionality](#). In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013.*, pages 3111–3119.
- Dimitrijs Milajevs, Dimitri Kartsaklis, Mehrnoosh Sadrzadeh, and Matthew Purver. 2014. Evaluating neural word representations in tensor-based compositional settings. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 708–719.
- George A. Miller. 1995. [Wordnet: A lexical database for english](#). *Commun. ACM*, 38(11):39–41.
- George A. Miller and Walter G. Charles. 1991. [Contextual correlates of semantic similarity](#). *Language and Cognitive Processes*, 6(1):1–28.
- Jeff Mitchell and Mirella Lapata. 2008. Vector-based models of semantic composition. In *Proceedings of ACL-08: HLT*, pages 236–244.
- Roberto Navigli. 2009. [Word sense disambiguation: A survey](#). *ACM Comput. Surv.*, 41(2).
- Arvind Neelakantan, Jeevan Shankar, Alexandre Passos, and Andrew McCallum. 2014. [Efficient non-parametric estimation of multiple embeddings per word in vector space](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1059–1069, Doha, Qatar. Association for Computational Linguistics.
- Denis Paperno, Nghia The Pham, and Marco Baroni. 2014. [A practical and linguistically-motivated approach to compositional distributional semantics](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 90–99, Baltimore, Maryland. Association for Computational Linguistics.
- Robin Piedeleu, Dimitri Kartsaklis, Bob Coecke, and Mehrnoosh Sadrzadeh. 2015. Open system categorical quantum semantics in natural language processing. In *Proceedings of the 6th Conference on Algebra and Coalgebra in Computer Science*, volume 35, pages 270–289. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.
- Mohammad Taher Pilehvar and José Camacho-Collados. 2019. WiC: the word-in-context dataset for evaluating context-sensitive meaning representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1267–1273.
- Alessandro Raganato, Jose Camacho-Collados, and Roberto Navigli. 2017. [Word sense disambiguation: A unified evaluation framework and empirical comparison](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, pages 99–110, Valencia, Spain. Association for Computational Linguistics.
- Herbert Rubenstein and John B. Goodenough. 1965. [Contextual correlates of synonymy](#). *Commun. ACM*, 8(10):627–633.
- Mehrnoosh Sadrzadeh, Dimitri Kartsaklis, and Esmá Balkir. 2018. [Sentence entailment in compositional distributional semantics](#). *Annals of Mathematics and Artificial Intelligence*, 82(4):189–218.

- Mehrnoosh Sadrzadeh and Reinhard Muskens. 2018. Static and dynamic vector semantics for lambda calculus models of natural language. *Journal of Language Modelling* Vol. 6(2):319–351.
- Hinrich Schütze. 1998. Automatic word sense discrimination. *Computational Linguistics*, 24(1):97–123.
- Richard Socher, Brody Huval, Christopher D Manning, and Andrew Y Ng. 2012. Semantic compositionality through recursive matrix-vector spaces. In *Proceedings of the 2012 joint conference on empirical methods in natural language processing and computational natural language learning*, pages 1201–1211.
- Gregor Wiedemann, Steffen Remus, Avi Chawla, and Chris Biemann. 2019. Does BERT make any sense? Interpretable word sense disambiguation with contextualized embeddings. In *Proceedings of KONVENS 2019*.
- Gijs Wijnholds and Mehrnoosh Sadrzadeh. 2019. [Evaluating composition models for verb phrase elliptical sentence embeddings](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 261–271, Minneapolis, Minnesota. Association for Computational Linguistics.

A Word2DM Gradients

The objective function that SGNS optimises at each prediction with regard to model parameters θ is

$$J(\theta) = \log \sigma(v_t^\top v_c) + \sum_{k=1}^K \log \sigma(-v_t^\top v_{w_k}) \quad (6)$$

where v_t is the embedding of target word, v_c is the embedding of the context word, and v_1, v_2, \dots, v_K are the embeddings of K negative samples. By optimising equation 6 over a large corpus, skip-gram learns word embeddings that encode distributional information.

Maximising equation 6 adjusts the embeddings of words occurring in the same context to be more similar and adjusts the embeddings of words that don't occur together to be less similar. This becomes clear when we consider the gradients used to update embeddings during training. We briefly recall the details of the gradient calculation so as to refer back to it later in this section. The derivative of equation 6 with respect to the target vector v_t is

$$\frac{\partial J}{\partial v_t} = (1 - \sigma(v_t^\top v_c))v_c - \sum_{k=1}^K (1 - \sigma(v_t^\top v_{w_k}))v_{w_k} \quad (7)$$

which is used to update the target vector as follows:

$$v_t \leftarrow v_t + \alpha \frac{\partial J}{\partial v_t} \quad (8)$$

The target vector is updated by adding the scaled context vector to it and subtracting the scaled negatively sampled vectors from it. The vectors are scaled proportionally to how dissimilar they are to the target vector. This ensures that the target vector is ‘‘pulled closer’’ to the true context vector and ‘‘pushed away’’ from the negative context vectors. It is this computationally simple training procedure which makes SGNS effective.

Word2DM extends SGNS to learn density matrices, replacing equation 6 with the following objective function:

$$J(\theta) = \log \sigma(\text{tr}(A_t A_c)) + \sum_{k=1}^K \log \sigma(-\text{tr}(A_t A_{w_k})) \quad (9)$$

where A_t and A_c are the density matrices of the target and context words respectively, A_1, A_2, \dots, A_K are the density matrices of K negative samples, and θ is the set of weights of the intermediary matrices B_t, B_c and B_1, B_2, \dots, B_K .

Computing this objective function requires multiple matrix multiplications. For each $\text{tr}(A_t A_c)$ term (including the terms of the K negative samples), the matrices A_t and A_c have to be computed respectively as $A_t = B_t B_t^\top$ and $A_c = B_c B_c^\top$ and then the matrix product $A_t A_c$ has to be computed. This means that, for each target-context prediction, we require $3(K+1)$ matrix multiplications. One of the most attractive features of SGNS is its computational efficiency, which enabled training on very large corpora in reasonable time. The introduction of multiple matrix multiplications into the objective function means that much of this efficiency is lost. In order to reduce the complexity of our model, we make use of the following property and lemma to find a new objective function that is computationally simpler, but equivalent to equation 9.

Property A.1. *The trace of the product of two matrices can be expressed as the sum of the element-wise products of their elements. If A is an $n \times m$ matrix and B is an $m \times n$ matrix, then the trace of the $n \times n$ matrix AB can be computed as*

$$\text{tr}(AB) = \sum_{i=1}^n \sum_{j=1}^m a_{ij} b_{ji}$$

Lemma A.2. *If B_t and B_c are $n \times m$ intermediary matrices, then trace of the matrix product $A_t A_c$ can be written as the sum of the squared elements of an $m \times m$ matrix $C = B_c^\top B_t$:*

$$\text{tr}(A_t A_c) = \sum_{i=1}^m \sum_{j=1}^m c_{ij}^2$$

Proof. We can express $\text{tr}(A_t A_c)$ as a trace computation involving intermediary matrices B_t and B_c :

$$\text{tr}(A_t A_c) = \text{tr}(B_t B_t^\top B_c B_c^\top)$$

Then we can use the cyclic property of the trace function to rewrite this as the product of a matrix C and its transpose:

$$\begin{aligned} \text{tr}(A_t A_c) &= \text{tr}(B_c^\top B_t B_t^\top B_c) \\ &= \text{tr}(B_c^\top B_t (B_c^\top B_t)^\top) \\ &= \text{tr}(C C^\top), \quad \text{where } C = B_c^\top B_t \end{aligned}$$

Now we can use property A.1 to express this as the element-wise products of the elements of C and its

transpose:

$$\begin{aligned}\text{tr}(A_t A_c) &= \sum_{i=1}^m \sum_{j=1}^m [C]_{ij} [C^\top]_{ji} \\ &= \sum_{i=1}^m \sum_{j=1}^m c_{ij} c_{ij} \\ &= \sum_{i=1}^m \sum_{j=1}^m c_{ij}^2\end{aligned}$$

□

This allows us to rewrite equation 9 to find an equivalent objective function that requires fewer computations than straightforward matrix multiplication would. The objective function at each target-context prediction becomes

$$\begin{aligned}J(\theta) &= \log \sigma \left(\sum_{i=1}^m \sum_{j=1}^m [B_c^\top B_t]_{ij}^2 \right) \\ &+ \sum_{k=1}^K \log \sigma \left(- \sum_{i=1}^m \sum_{j=1}^m [B_{w_k}^\top B_t]_{ij}^2 \right).\end{aligned}\quad (10)$$

By using the result of lemma A.2 we have reduced the number of matrix multiplications required for each target-context prediction from $3(K+1)$ to $(K+1)$. Density matrices are trained by maximising equation 10 with respect to the intermediary matrices $B_t, B_c, B_{w_1}, \dots, B_{w_K}$ over a large corpus.

The model is trained using stochastic gradient descent. We now derive the gradients used to update B_t during training, and subsequently show that these gradients lead to suboptimal updates to the density matrices during training. Deriving the gradient with respect to B_c and B_{w_k} would proceed similarly. To compute the gradients of equation 10 we first rewrite it in terms of the elements of the $n \times m$ matrices B_t, B_c , and B_{w_k} :

$$\begin{aligned}J(\theta) &= \log \sigma \left(\sum_{i=1}^m \sum_{j=1}^m \left(\sum_{l=1}^n b_{li}^c b_{lj}^t \right)^2 \right) \\ &+ \sum_{k=1}^K \log \sigma \left(- \sum_{i=1}^m \sum_{j=1}^m \left(\sum_{l=1}^n b_{li}^{w_k} b_{lj}^t \right)^2 \right),\end{aligned}\quad (11)$$

where b_{pq}^x denotes the pq th element of B_x . We derive the gradient of this objective function with respect to b_{pq}^t , an element of the intermediary target word matrix B_t . In order to use the chain rule in

gradient calculations we rewrite $J(\theta)$ as a composite function:

$$J(\theta) = \log \sigma(y(\theta)) + \sum_{k=1}^K \log \sigma(z_k(\theta)), \quad (12)$$

where

$$y(\theta) = \sum_{i=1}^m \sum_{j=1}^m \left(\sum_{l=1}^n b_{li}^c b_{lj}^t \right)^2 \quad \text{and}$$

$$z_k(\theta) = - \sum_{i=1}^m \sum_{j=1}^m \left(\sum_{l=1}^n b_{li}^{w_k} b_{lj}^t \right)^2.$$

The derivative of J with respect to b_{pq}^t can now be computed as follows:

$$\begin{aligned}\frac{\partial J}{\partial b_{pq}^t} &= \frac{\partial \log \sigma}{\partial \sigma} \frac{\partial \sigma}{\partial y} \frac{\partial y}{\partial b_{pq}^t} + \sum_{k=1}^K \frac{\partial \log \sigma}{\partial \sigma} \frac{\partial \sigma}{\partial z_k} \frac{\partial z_k}{\partial b_{pq}^t} \\ &= \frac{1}{\sigma(y)} (1 - \sigma(y)) \sigma(y) \frac{\partial y}{\partial b_{pq}^t} \\ &+ \sum_{k=1}^K \frac{1}{\sigma(z_k)} (1 - \sigma(z_k)) \sigma(z_k) \frac{\partial z_k}{\partial b_{pq}^t} \\ &= (1 - \sigma(y)) \frac{\partial y}{\partial b_{pq}^t} + \sum_{k=1}^K (1 - \sigma(z_k)) \frac{\partial z_k}{\partial b_{pq}^t} \\ &= (1 - \sigma(y)) \sum_{i=1}^m \sum_{l=1}^n 2b_{pi}^c \sum_{l=1}^n b_{li}^c b_{lj}^t \\ &- \sum_{k=1}^K (1 - \sigma(z_k)) \sum_{i=1}^m 2b_{pi}^{w_k} \sum_{l=1}^n b_{li}^{w_k} b_{lj}^t \\ &= (1 - \sigma(y)) 2[B_c B_c^\top B_t]_{pq} \\ &- \sum_{k=1}^K (1 - \sigma(z_k)) 2[B_{w_k} B_{w_k}^\top B_t]_{pq}\end{aligned}$$

The last line in the above derivation is obtained by rewriting the summation expressions as equivalent matrix multiplications. We can now write the derivative of J with respect to the full intermediary matrix B_t :

$$\begin{aligned}\frac{\partial J}{\partial B_t} &= (1 - \sigma(y(\theta))) 2B_c B_c^\top B_t \\ &+ \sum_{k=1}^K (1 - \sigma(z_k(\theta))) 2B_{w_k} B_{w_k}^\top B_t\end{aligned}\quad (13)$$

As opposed to the gradients of Word2Vec (equation 7), the gradients of Word2DM do not lead to simple and easily interpretable training updates. As discussed in the paragraph following equation 8, in Word2Vec the target vector is made more similar to the context vector and less similar to the negative context vectors. Ideally we would like something

similar to occur in Word2DM with density matrices, but equation 13 shows that we lose the intuitive training updates of Word2Vec through the introduction of intermediary matrices. Furthermore, we can show that the gradients of Word2DM sometimes lead to unwanted consequences in training.

Consider the case where the density matrices of a target and context word are highly dissimilar. Recall from equation 9 that the y is in equation 13 is the trace inner product of the density matrices A_t and A_c (the measure we use to quantify semantic similarity). The minimum value of the trace inner product of two density matrices is zero (this follows from the fact that density matrices are positive semi-definite), so two density matrices are highly dissimilar when their trace inner product is close to zero i.e. $y \approx 0$. From equation 10 we can recall how y can be written in terms of the intermediary matrices:

$$y = \sum_{i=1}^m \sum_{j=1}^m [B_c^\top B_t]_{ij}^2$$

Consider that $y \approx 0$ if and only if the elements of $B_c^\top B_t$ are close to zero in value, since squaring the elements in the summation makes them all positive. We have established the following equivalence:

$$\text{tr}(A_t A_c) \approx 0 \iff B_c^\top B_t \approx O,$$

where O is the $m \times m$ matrix with all zero entries. Consider how this will affect the target-context update during training. The first term of the gradient in equation 13 becomes

$$(1 - \sigma(y(\theta))) 2B_c B_c^\top B_t = (1 - \sigma(0)) 2B_c O \approx O$$

so the target-context update becomes ineffective for true contexts. The update should make the density matrix of the target word more similar to that of the context word, but the gradient is so small that it makes this impossible. Moreover, the more dissimilar the target and context density matrices are before the update, the less effective the update will be. This is the opposite of the intended effect (achieved by Word2Vec) in which the magnitude of the target-context update should increase if the target and context representations are less similar. This is an example of how the introduction of intermediary matrices in Word2DM leads to suboptimal training updates. We ensure that our density matrices are positive semi-definite, but lose the guarantee that the algorithm will learn high-quality semantic representations.

B Hyperparameters for Word2DM and Multi-Sense Word2DM

Word2DM additional details We use a dynamic window size i.e. the size of each context window is sampled uniformly between 1 and the maximum window size. We also discard words that occur less than some minimum threshold and subsample frequently occurring words. Negative samples are drawn from a unigram distribution raised to the power of $\frac{3}{4}$. Furthermore, we train two density matrices for each word - one that represents it as a target word and another that represents it as a context word. After training we use the target density matrices as our final density matrices.

Hyperparameters We train our Word2DM and multi-sense Word2DM models on the ukWaC+Wackypedia corpus, consisting of 2.8 billion words. We use a window size of 5, a minimum word count of 50, 5 negative samples per positive context, and a subsampling rate of 1e-5. We train the model for 4 iterations of the ukWaC+Wackypedia corpus, using the Adam optimisation algorithm, a learning rate of 0.001, and 16 sentences per batch.