

Density Matching for Bilingual Word Embedding

Chunting Zhou, Xuezhe Ma, Di Wang, Graham Neubig

Language Technologies Institute

Carnegie Mellon University

ctzhou, xuezhem, diwang, gneubig@cs.cmu.edu

Abstract

Recent approaches to cross-lingual word embedding have generally been based on linear transformations between the sets of embedding vectors in the two languages. In this paper, we propose an approach that instead expresses the two monolingual embedding spaces as probability densities defined by a Gaussian mixture model, and matches the two densities using a method called normalizing flow. The method requires no explicit supervision, and can be learned with only a seed dictionary of words that have identical strings. We argue that this formulation has several intuitively attractive properties, particularly with the respect to improving robustness and generalization to mappings between difficult language pairs or word pairs. On a benchmark data set of bilingual lexicon induction and cross-lingual word similarity, our approach can achieve competitive or superior performance compared to state-of-the-art published results, with particularly strong results being found on etymologically distant and/or morphologically rich languages.¹

1 Introduction

Cross-lingual word embeddings represent words in different languages in a single vector space, capturing the syntactic and semantic similarity of words across languages in a way conducive to use in computational models (Upadhyay et al., 2016; Ruder et al., 2017). These embeddings have been shown to be an effective tool for cross-lingual NLP, e.g. the transfer of models trained on high-resource languages to low-resource ones (Klementiev et al., 2012; Guo et al., 2015; Zoph et al., 2016; Zhang et al., 2018; Gu et al., 2018) or unsupervised learning (Artetxe et al., 2018c).

¹Code/scripts can be found at <https://github.com/violet-zct/DeMa-BWE>.

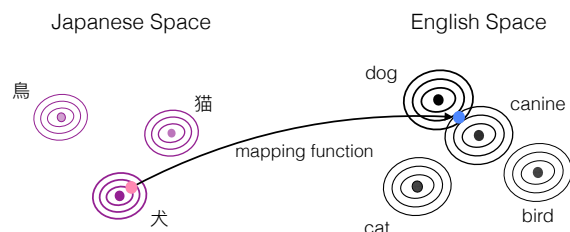


Figure 1: An illustration of our method. Thicker lines represent higher mixture weights, linked to word frequency. The pink point is a continuous training sample in the Japanese embedding space while the blue point is a mapped point in the English space.

There are two major paradigms in the learning of cross-lingual word embeddings: “online” and “offline”. “Online” methods learn the cross-lingual embeddings directly from parallel corpora (Hermann and Blunsom, 2014), optionally augmented with monolingual corpora (Gouws et al., 2015). In contrast, “offline” approaches learn a bilingual mapping function or multilingual projections from pre-trained monolingual word embeddings or feature vectors (Haghighi et al., 2008; Mikolov et al., 2013; Faruqui and Dyer, 2014). In this work, we focus on this latter offline approach.

The goal of bilingual embedding is to learn a shared embedding space where words possessing similar meanings are projected to nearby points. Early work focused on supervised methods maximizes the similarity of the embeddings of words that exist in a manually-created dictionary, according to some similarity metric (Mikolov et al., 2013; Faruqui and Dyer, 2014; Jawanpuria et al., 2018; Joulin et al., 2018). In contrast, recently proposed unsupervised methods frame this problem as minimization of some form of distance between the whole set of discrete word vectors in the chosen vocabulary, e.g. Wasserstein distance or Jensen–Shannon divergence (Xu et al., 2018; Conneau et al., 2017; Zhang et al., 2017; Grave

et al., 2018). While these methods have shown impressive results for some language pairs despite the lack of supervision, regarding the embedding space as a set of discrete points has some limitations. First, expressing embeddings as a single point in the space doesn't take into account the inherent uncertainty involved in learning embeddings, which can cause embedding spaces to differ significantly between training runs (Wendlandt et al., 2018). Second, even in a fixed embedding space the points surrounding those of words that actually exist in the pre-trained vocabulary also often are coherent points in the embedding space.

In this work, we propose a method for **density matching for bilingual word embedding (DeMa-BWE)**. Instead of treating the embedding space as a collection of discrete points, we express it as a probability density function over the entire continuous space over word vectors. We assume each vector in the monolingual embedding space is generated from a Gaussian mixture model with components centered at the pretrained word embeddings (Fig. 1), and our approach then learns a bilingual mapping that most effectively matches the two probability densities of the two monolingual embedding spaces.

To learn in this paradigm, instead of using the pre-trained word embeddings as fixed training samples, at every training step we obtain samples from the Gaussian mixture space. Thus, our method is exploring the entire embedding space instead of only the specific points assigned for observed words. To calculate the density of the transformed samples, we use volume-preserving invertible transformations over the target word embeddings, which make it possible to perform density matching in a principled and efficient way (Rezende and Mohamed, 2015; Papamakarios et al., 2017; He et al., 2018). We also have three additional ingredients in the model that proved useful in stabilizing training: (1) a back-translation loss to allow the model to learn the mapping jointly in both directions, (2) an identical-word-matching loss that provides weak supervision by encouraging the model to have words with identical spellings be mapped to a similar place in the space, and (3) frequency-matching based Gaussian mixture weights that accounts for the approximate frequencies of aligned words.

Empirical results are strong; our method is able to effectively learn bilingual embeddings

that achieve competitive or superior results on the MUSE dataset (Conneau et al., 2017) over state-of-the-art published results on bilingual word translation and cross-lingual word similarity tasks. The results are particularly encouraging on etymologically distant or morphologically rich languages, as our model is able to explore the integration over the embedding space by treating the space as a continuous one. Moreover, unlike previous unsupervised methods that are usually sensitive to initialization or require sophisticated optimization procedures, our method is robust and requires no special initialization.

2 Background: Normalizing Flows

In this section, we will briefly describe normalizing flows - the backbone of DeMa-BWE.

As mentioned in the introduction and detailed later, our model is based on matching two probability density functions, one representing the source embedding space and one representing the target embedding space. To learn in this framework, we will use the concept of *normalizing flows* (Rezende and Mohamed, 2015). We will explain them briefly here, but refer readers to Rezende and Mohamed (2015) for details due to space constraints.

Concretely, let \mathbf{u} denote a high dimensional random vector (e.g. representing a point in the source embedding space) and \mathbf{z} be a latent variable that corresponds to \mathbf{u} (e.g. a point in the target embedding space). Flow-based generative models (Kingma et al., 2016; Kingma and Dhariwal, 2018) learn invertible transformations $f_\theta(\mathbf{z})$ from the distribution over \mathbf{z} to the distribution over \mathbf{u} . The generative story of the model is defined as:

$$\mathbf{z} \sim p_\theta(\mathbf{z}), \quad \mathbf{u} = f_\theta(\mathbf{z}) \quad (1)$$

where $p_\theta(\mathbf{z})$ is the prior distribution. This prior can be any distribution for which we can tractably compute the density of sample points \mathbf{z} . A common choice of such distribution is a spherical multivariate Gaussian distribution: $p_\theta(\mathbf{z}) = \mathcal{N}(\mathbf{z}; 0, \mathbf{I})$ (Dinh et al., 2016). Assuming the transformation function $f_\theta(\cdot)$ is invertible, using the rule for change of variables, the probability density of \mathbf{u} can be calculated as:

$$p_\theta(\mathbf{u}) = p_\theta(\mathbf{z})|\det(J(f^{-1}(\mathbf{u})))| \quad (2)$$

where $\det(J(f^{-1}(\mathbf{u})))$ is determinant of the Jacobian matrix of the function inverse. This term

accounts for the way in which f locally expands or contracts regions of \mathbf{z} , and enforces the invertibility of the function. A “normalizing flow” is a cascaded sequence of such invertible transformations, which is learned by maximizing the density in Equation (2) over observed data points \mathbf{u} . One computational issue with these models lies in calculating the Jacobian matrix, which is expensive in the general case. A common method is to choose transformations whose Jacobians’ are a triangular matrix, which renders this computation tractable (Dinh et al., 2016; Papamakarios et al., 2017).

3 Proposed Method

In this section, we present notation used in our method, describe the prior we define for the monolingual embedding space, then detail our density matching method.

3.1 Notation

Given two sets of independently trained monolingual embeddings, the problem of bilingual embedding mapping is to learn a mapping function that aligns the two sets in a shared space. Let $\mathbf{x} \in \mathbb{R}^d, \mathbf{y} \in \mathbb{R}^d$ denote vectors in the source and target language embedding space respectively. Let x_i and y_j denote an actual word in the pretrained source and target vocabularies respectively. Words are sorted by their occurrence counts in the monolingual corpus and the index of the word represents its rank. We use \mathbf{x}_i and \mathbf{y}_j to denote the pretrained word embeddings for word x_i in the source language and word y_j in the target language respectively. Given a pair of languages s and t , our approach learns two mapping functions: f_{xy} that maps the source embedding to the target space and f_{yx} that gives the mapping in the reverse direction.

3.2 Density Estimation in Monolingual Space

To learn the mapping, we project a vector \mathbf{x} in the source embedding space into the target space \mathbf{y} . We learn this mapping by maximizing the density of data points in the source space. The density can be computed using the idea of normalizing flow described above. Thus, for the the monolingual embedding spaces, we need to define tractable density functions $p(\mathbf{x})$ and $p(\mathbf{y})$.

While any number of functions could be conceived to calculate these densities, in the current method we opt to use a Gaussian Mixture Model (GMM) with Gaussian components cen-

tered at each pretrained word embedding. This is motivated by the assumption that embeddings are likely to appear in the neighborhood of other embeddings, where we define the “neighborhood” to be characterized as closeness in Euclidean space, and the uncertainty of each neighborhood as being Gaussian. Concretely, let N_x and N_y denote the number of pretrained word embeddings that serve as Gaussian component centers during training for the source and target languages, respectively. Then we can express the density of any point in the source embedding space as:

$$p(\mathbf{x}) = \sum_{i \in \{1, \dots, N_x\}} \pi(x_i) \tilde{p}(\mathbf{x}|x_i) \quad (3)$$

where $\pi(x_i)$ is the frequency of word x_i normalized within the N_x component words, and $\tilde{p}(\mathbf{x}|x_i)$ is a Gaussian distribution centered at the embedding of word x_i . We simply use a fixed variance σ_x^2 for all Gaussian components:

$$\tilde{p}(\mathbf{x}|x_i) = \mathcal{N}(\mathbf{x}|\mathbf{x}_i, \sigma_x^2 \mathbf{I}) \quad (4)$$

Similarly, the density of any point in the target embedding space can be written as:

$$p(\mathbf{y}) = \sum_{j \in \{1, \dots, N_y\}} \pi(y_j) \tilde{p}(\mathbf{y}|y_j) \quad (5)$$

where $\tilde{p}(\mathbf{y}|y_j) = \mathcal{N}(\mathbf{y}|\mathbf{y}_j, \sigma_y^2 \mathbf{I})$.

3.3 Density Matching

With the Gaussian mixture model as the prior distribution in the monolingual space, our goal is to learn a mapping function from one embedding space to the other such that the log probabilistic density is maximized in the source space.

While we are jointly learning the two mapping functions f_{xy} and f_{yx} simultaneously, for conciseness we will illustrate our approach using the source to target mapping f_{xy} . First, a continuous vector \mathbf{x} is sampled from the Gaussian mixture model (Eq. (3)) by sampling $x_i \sim \pi(x_i)$ then $\mathbf{x} \sim \tilde{p}(\mathbf{x}|x_i)$ (4). Next, we apply the mapping function f_{xy} to obtain the transformed vector \mathbf{y} in the target space. Concretely, the mapping functions we employ in this work are two linear transformations: $f_{xy}(\cdot) = \mathbf{W}_{xy} \cdot$ and $f_{yx} = \mathbf{W}_{yx} \cdot$. Connecting to the transformation function in Sec. 2, we see that $\mathbf{x} = f(\mathbf{y}) = \mathbf{W}_{xy}^{-1} \mathbf{y}$, $\mathbf{y} = f^{-1}(\mathbf{x}) = \mathbf{W}_{xy} \mathbf{x}$, and $J(f^{-1}(\mathbf{x})) = \mathbf{W}_{xy}$. We can then express the log density of a sample \mathbf{x} as:

$$\log p(\mathbf{x}; \mathbf{W}_{xy}) = \log p(\mathbf{y}) + \log |\det(\mathbf{W}_{xy})| \quad (6)$$

where the Jacobian regularization term accounts for the volume expansion or contraction resulting from the projection matrix \mathbf{W}_{xy} .

We maximize the likelihood function which is equivalent to minimizing expectation of the KL-divergence between the prior distribution and the model distribution of \mathbf{x} . This provides a natural objective for optimization:

$$\text{minimize: } \text{KL}(p(\mathbf{x})||p(\mathbf{x}; \mathbf{W}_{xy})) \quad (7)$$

By replacing $\mathbf{W}_{xy}\mathbf{x}$ with \mathbf{y} , this is equivalent to maximizing the log density of transformed source samples in the target space (see Eq. (6)):

$$\mathcal{L}_{xy} = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log p(\mathbf{y}) + \log |\det(\mathbf{W}_{xy})|] \quad (8)$$

The objective \mathcal{L}_{xy} contains two parts: the log density function $\log p(\mathbf{y})$ and a regularization term $\log \det(\mathbf{W}_{xy})$. Likewise, for the target to source mapping \mathbf{W}_{yx} , we have the density matching objective \mathcal{L}_{yx} .

Conditional Density Matching The above marginal density matching method does not take into account the dependency between the embeddings in the two monolingual spaces. To address this issue, we extend the density matching method to the conditional density function:

$$\log p(\mathbf{x}|x_i; \mathbf{W}_{xy}) = \log p(\mathbf{y}|x_i) + \log |\det(\mathbf{W}_{xy})|$$

The conditional density $p(\mathbf{y}|x_i)$ is the prior distribution in this simple normalizing flow, and for this we use a Gaussian mixture model in the target monolingual space:

$$\begin{aligned} p(\mathbf{y}|x_i) &= \sum_{j \in \{1, \dots, N_y\}} p(\mathbf{y}, y_j|x_i) \\ &= \sum_{j \in \{1, \dots, N_y\}} \tilde{p}(\mathbf{y}|y_j)\pi(y_j|x_i) \quad (9) \end{aligned}$$

Similarly, where $\tilde{p}(\mathbf{y}|y_j)$ is the Gaussian density function in the mixture model defined in Equation (5). $\pi(y_j|x_i)$ allows us to incorporate a-priori assumptions about whether two words are likely to match. In fact, the density matching method in Eq. (6) can be regarded as a special case of the conditional density matching method by adopting a naive prior $\pi(y_j|x_i) := \pi(y_j)$. However previous work (Zhang et al., 2017) has noted that word frequency information is a strong signal – words with similar frequency are likely to be matched –

and thus we use $\pi(y_j|x_i)$ to incorporate this prior knowledge.

In this work, we assume that the frequencies of aligned bilingual words in the individual monolingual corpus should be correlated, and to match words that are ranked similarly, we model the Gaussian mixture weights as the negative absolute difference between log-scale word ranks and normalize over all the target Gaussian component words by a softmax function with temperature τ :

$$\pi(y_j|x_i) = \frac{\exp(-|\log(j) - \log(i)|/\tau)}{\sum_{k=1}^{N_y} \exp(-|\log(k) - \log(i)|/\tau)}$$

Thus, if a word x_i has similarly frequency rank as word y_j , the sample \mathbf{x} from the Gaussian distribution centered at \mathbf{x}_i will be assigned higher weight for the component $\tilde{p}(\mathbf{y}|y_j) = \tilde{p}(\mathbf{W}_{xy}\mathbf{x}|y_j)$. Although this assumption will not hold always (e.g. for languages that have different levels of morphological complexity), intuitively we expect that using this signal will help more overall than it will hurt, and empirically we find that this weighting is not sensitive to language variation and works well in practice.

The updated objective is

$$\begin{aligned} \mathcal{L}_{xy} &= \mathbb{E}_{x_i \sim \pi(x_i)} [\text{KL}(\tilde{p}(\mathbf{x}|x_i)||p(\mathbf{x}|x_i; \mathbf{W}_{xy}))] \\ &= \mathbb{E}_{x_i \sim \pi(x_i), \mathbf{x} \sim \tilde{p}(\mathbf{x}|x_i)} [\log p(\mathbf{y}|x_i) \\ &\quad + \log |\det(\mathbf{W}_{xy})|] \quad (10) \end{aligned}$$

In the conditional density above, both the frequency-matching weight and the Gaussian density function play an important role in matching the density of a source-space sample with the target embedding space. The former matches bilingual words with their frequency ranks while the latter matches words with their vector distances.

3.4 Weak Orthogonality Constraint

A common choice of bilingual mapping function is a linear transformation with an orthogonality constraint. Various motivations have been proposed for the orthogonality constraint such as length normalization of embeddings (Xing et al., 2015), and reversible mapping (Smith et al., 2017). In this work, we add a weak orthogonality constraint to the bilingual mappings via a back-translation loss as follows:

$$\begin{aligned} \mathcal{L}_{bt} &= \mathbb{E}_{x_i \sim \pi(x_i), \mathbf{x} \sim \tilde{p}(\mathbf{x}|x_i)} [g(\mathbf{W}_{yx}\mathbf{W}_{xy}\mathbf{x}, \mathbf{x})] \\ &\quad + \mathbb{E}_{y_j \sim \pi(y_j), \mathbf{y} \sim \tilde{p}(\mathbf{y}|y_j)} [g(\mathbf{W}_{xy}\mathbf{W}_{yx}\mathbf{y}, \mathbf{y})] \end{aligned}$$

where $g(\cdot, \cdot) = 1 - \text{cosine}(\cdot, \cdot)$ is the cosine loss. Jointly learning the two mapping matrices by minimizing this cyclic loss encourages \mathbf{W}_{xy} and \mathbf{W}_{yx} to be orthogonal to each other.

3.5 Weak Supervision with Identical Words

To reduce the search space of the mapped bilingual embeddings, we add an additional weakly supervised loss over words that have identical strings in both the source and target languages denoted \mathcal{W}_{id} .

$$\mathcal{L}_{sup} = \sum_{v \in \mathcal{W}_{id}} g(\mathbf{v}_x \mathbf{W}_{xy}^T, \mathbf{v}_y) + g(\mathbf{v}_y \mathbf{W}_{yx}^T, \mathbf{v}_x)$$

where \mathbf{v}_x and \mathbf{v}_y are the pretrained word embedding of word v in the source and target side respectively, and $g(\cdot, \cdot)$ is the cosine loss described above. Although using identical strings for supervision is very noisy, especially for languages with little overlap in vocabularies, we find that they provide a enough guidance to training to prevent the model from being trapped in poor local optima.

Putting everything together, the overall objective function of DeMa-BWE includes three parts: the density matching loss, the weak orthogonality loss and the weak supervised loss:

$$\mathcal{L} = \mathcal{L}_{xy} + \mathcal{L}_{yx} + \lambda \cdot \mathcal{L}_{bt} + \alpha \cdot \mathcal{L}_{sup} \quad (11)$$

where λ and α are coefficients that tradeoff between different losses.

4 Retrieval and Refinement

4.1 Retrieval Method

One standard use case for bilingual embeddings is in bilingual lexicon induction, where the embeddings are used to select the most likely translation in the other language given these embeddings. In this case, it is necessary to have a retrieval metric that selects word or words likely to be translations given these embeddings. When performing this retrieval, it has been noted that high-dimensional embedding spaces tend to suffer from the ‘‘hubness’’ problem (Radovanović et al., 2010) where some vectors (known as *hubs*) are nearest neighbors of many other points, which is detrimental to reliably retrieving translations in the bilingual space. To mitigate the hubness problem, we adopt the *Cross-Domain Similarity Local Scaling* (CSLS) metric proposed in (Conneau et al., 2017) that penalizes the similarity score of

the hubs. Specifically, given two mapped embeddings $\mathbf{W}_{xy}\mathbf{x}$ denoted \mathbf{x}' and \mathbf{y} , CSLS first computes the average cosine similarity of \mathbf{x}' and \mathbf{y} for their k nearest neighbors denoted $r_T(\mathbf{x}')$ and $r_S(\mathbf{y})$ in the other language respectively, then the corrected similarity measure $\text{CSLS}(\cdot, \cdot)$ is defined as:

$$\text{CSLS}(\mathbf{x}', \mathbf{y}) = 2\text{cos}(\mathbf{x}', \mathbf{y}) - r_T(\mathbf{x}') - r_S(\mathbf{y})$$

where $\text{cos}(\cdot, \cdot)$ is the cosine similarity. Following (Conneau et al., 2017), k is set to be 10.

CSLS consistently outperform cosine similarity on nearest neighbor retrieval, however it does not consider the relative frequency of bilingual words which we hypothesize can be useful in disambiguation. As we discussed in Sec. 3.2, our density matching objective considers both the relative frequencies and vector similarities. The conditional density $p(\mathbf{y}|x_i)$ (Eq. (9)) in our density matching objective (Eq. (8)) is a marginalized distribution over all target component words y_j where the density of each component $p(\mathbf{y}, y_j|x_i)$ can be directly used as a similarity score for a pair of words (y_j, x_i) to replace the cosine similarity $\text{cos}(\mathbf{x}', \mathbf{y})$ in CSLS. Let CSLS-D denote this modified CSLS metric, which we compare to standard CSLS in experiments. We find that using CSLS-D for nearest neighbor retrieval outperforms the CSLS metric in most cases on the bilingual dictionary induction task.

4.2 Iterative Procrustes Refinement

Iterative refinement, which learns the new mapping matrix by constructing a bilingual lexicon iteratively, has been shown as an effective method for improving the performance of unsupervised lexicon induction models (Conneau et al., 2017). Given a learned bilingual embedding mapping \mathbf{W} , the refinement starts by inferring an initial bilingual dictionary using the retrieval method above on the most frequent words. Let \mathbf{X} and \mathbf{Y} denote the ordered embedding matrices for the inferred dictionary words for source and target languages respectively. Then a new mapping matrix \mathbf{W}^* is induced by solving the Procrustes problem:

$$\begin{aligned} \mathbf{W}^* &= \underset{\mathbf{W} \in O_d(\mathbb{R})}{\text{argmin}} \|\mathbf{W}\mathbf{X} - \mathbf{Y}\|_F = \mathbf{U}\mathbf{V}^T \\ \text{s.t. } \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T &= \text{SVD}(\mathbf{Y}\mathbf{X}^T) \end{aligned}$$

The step above can be applied iteratively by inducing a new dictionary with the new mapping \mathbf{W} .

DeMa-BWE is able to achieve very competitive performance without further refinement, but for comparison we also report results with the refinement procedure, which brings small improvements in accuracy for most language pairs. Note that for bilingual dictionary induction during refinement, we use CSLS as the retrieval metric across all experiments for fair comparison to the refinement step in previous work.

5 Experiments

5.1 Dataset and Task

We evaluate our approach extensively on the bilingual lexicon induction (BLI) task, which measures the word translation accuracy in comparison to a gold standard. We report results on the widely used MUSE dataset (Conneau et al., 2017), which consists of FastText monolingual embeddings pre-trained on Wikipedia (Bojanowski et al., 2017), and dictionaries for many language pairs divided into train and test sets. We follow the evaluation setups of (Conneau et al., 2017). We evaluate DeMa-BWE by inducing lexicons between English and different languages including related languages, e.g. Spanish; etymologically distant languages, e.g. Japanese; and morphologically rich languages, e.g. Finnish.

5.2 Implementation Details

Embedding Normalization Following (Artetxe et al., 2018b), we pre-process the monolingual embeddings by first applying length normalization, then mean center each dimension, and then length normalize again to ensure that the final embeddings have a unit length. We observe that this normalization method helps stabilize training and accelerate convergence.

Other Experimental Details We held out 1000 translation pairs randomly sampled from the training set in the MUSE dataset as our validation data. We also tried the unsupervised validation criterion proposed in (Conneau et al., 2017) as the model selection method that computes the average cosine similarity over the model induced dictionary pairs and found that this unsupervised criterion can select models that achieve similar performance as the supervised validation criterion. All hyperparameters are tuned on the validation set and include the following: For the number of base words used as Gaussian components in the GMM, we

typically choose the most frequent 20,000 words for all language pairs but en-ja for which we use 10,000 which achieves better performance. We use a batch size of 2048 for all languages but en-ja for which we use 1024. We use Adam (Kingma and Ba, 2014) for optimization with default hyperparameters.

We empirically set the Gaussian variance to be 0.01 for both the source and target languages in en-es, en-de, en-fr, en-ru; in the experiments for morphologically rich languages (Sec. 5.4), we set the variance to be 0.015 for all these languages except for et whose variance is set to be 0.02 while keeping the variance of English to be 0.01. In the experiments for etymologically distant language pairs en-ja and en-zh, we set different variances for the source and target languages in different mapping directions. For details of the variance setting please check the scripts in our code base. We empirically find that for a language pair the variance of the language with relatively more complex morphological properties needed to be set larger than the other language, indicating that the model needs to explore more in the embedding space for the morphologically richer language.

We initialize mapping matrices \mathbf{W}_{xy} and \mathbf{W}_{yx} with a random orthogonal matrix. For the weak orthogonality constraint loss \mathcal{L}_{bt} , we set the weight λ to be 0.5 throughout all language pairs. For the weak supervision loss \mathcal{L}_{sup} , we set the weight α to be 10 for all languages except for en-zh where we find 5 performs better. We set the temperature τ used in the softmax function for Gaussian mixture weights to be 2 across all languages.

5.3 Main Results on BLI

In Tab. 1, we compare the performance of DeMa-BME extensively with the best performing unsupervised and supervised methods on the commonly benchmarked language pairs.

Our unsupervised baselines are: (1) **MUSE (U+R)** (Conneau et al., 2017), a GAN-based unsupervised method with refinement. (2) A strong and robust unsupervised self-learning method **SL-unsup** from (Artetxe et al., 2018b). We also run their published code with identical words as the initial dictionary for fair comparison with our approach, denoted **SL-unsup-ID**. (3) **Sinkhorn** (Xu et al., 2018) that minimizes the Sinkhorn distance between the source and target word vectors. (4) An iterative matching method from (Hoshen and

	en-es	es-en	en-de	de-en	en-fr	fr-en	en-ru	ru-en	en-zh	zh-en	en-ja	ja-en
Supervised												
Procrustes (R)	81.4	82.9	73.5	72.4	81.1	82.4	51.7	63.7	42.7	36.7	14.2	7.44
MSF-ISF	79.9	82.1	73.0	72.0	80.4	81.4	50.0	65.3	28.0	40.7	-	-
MSF	80.5	83.8	73.5	73.5	80.5	83.1	50.5	67.3	32.3	43.4	-	-
CSLS-Sp	80.7	83.9	75.1	72.1	81.7	83.2	51.1	63.8	-	-	-	-
GeoMM	81.4	85.5	74.7	76.7	82.1	84.1	51.3	67.6	49.1	45.3	-	-
Unsupervised												
MUSE (U+R)	81.7	83.3	74.0	72.2	82.3	81.1	44.0	59.1	32.5	31.4	0.0	4.2
SL-unsup	82.3	84.7	75.1	74.3	82.3	83.6	49.2	65.6	0.0	0.0	2.9	0.2
SL-unsup-ID	82.3	84.6	75.1	74.1	82.2	83.7	48.8	65.7	37.4	34.2	48.5	33.7
Sinkhorn*	79.5	77.8	69.3	67.0	77.9	75.5	-	-	-	-	-	-
Non-Adv	81.1	82.1	73.7	72.7	81.5	81.3	44.4	55.6	0.0	0.0	0.0	0.0
Non-Adv (R)	82.1	84.1	74.7	73.0	82.3	82.9	47.5	61.8	0.0	0.0	0.0	0.0
WS-Procrustes (R)	82.8	84.1	75.4	73.3	82.6	82.9	43.7	59.1	-	-	-	-
DeMa-BME												
CSLS (w/o R)	82.0	85.4	75.3	74.9	82.6	82.4	46.9	62.4	39.6	40.0	46.7	32.9
CSLS-D (w/o R)	82.3	85.1	76.3	75.1	83.7	82.5	48.0	61.7	40.5	37.7	45.3	32.4
CSLS (w/ R)	82.8	84.5	75.6	74.1	82.5	83.3	47.3	63.5	41.9	37.7	50.7	35.2
CSLS-D (w/ R)	82.8	84.9	77.2	74.4	83.1	83.5	49.2	63.6	42.5	37.9	52.0	35.6

Table 1: Precision@1 for the MUSE BLI task compared with previous work. All the baseline results employ CSLS as the retrieval metric except for Sinkhorn* which uses cosine similarity. **R** represents refinement. Bold and italic indicate the best unsupervised and overall numbers respectively. ('en' is English, 'es' is Spanish, 'de' is German, 'fr' is French, 'ru' is Russian, 'zh' is traditional Chinese, 'ja' is Japanese.)

Wolf, 2018)): **Non-Adv** and **Non-Adv (R)** with refinement. (5) **WS-Procrustes (R)** using refinement by (Grave et al., 2018). Our supervised methods include: (1) The iterative Procrustes method **Procrustes (R)** (Smith et al., 2017). (2) A multi-step framework **MSF-ISF** (Artetxe et al., 2018a) and its variant **MSF** which uses CSLS for retrieval, whose results are from (Jawanpuria et al., 2018). (3) **CSLS-Sp** by (Joulin et al., 2018) that optimizes the CSLS score, and (4) a geometric approach **GeoMM** by (Jawanpuria et al., 2018). For fair comparisons, all supervised results are trained with the training dictionaries in the MUSE dataset. All baseline methods employ CSLS for retrieval except for the Sinkhorn method.

For DeMa-BME, we present results with and without refinement, and with CSLS and CSLS-D as retrieval methods. From Tab. 1, we can see the overall performance of DeMa-BME is remarkable comparing with other unsupervised methods and is also competitive with strong supervised methods. The results without the iterative refinement **CSLS (w/o R)** are strong on almost all language pairs with particularly strong performance being observed on es-en, en-de and en-fr on which

DeMa-BME outperforms or is on par with the best performing methods. Applying refinement to DeMa-BME brings slightly better performance on most language pairs but degrades the performance on some language pairs such as es-en, zh-en for which the DeMa-BME already obtains very good results.

Also, DeMa-BME demonstrates notably better performance on distant language pairs (en-ru, en-ja and en-zh) over other unsupervised methods, which often achieve good performance on etymologically close languages but fail to converge on the distant language pairs. However, when the dictionary is initialized with identical strings for SL-unsup, we obtain decent results on these languages. The strong performance of supervised methods on Russian and Chinese demonstrates that on some language pairs supervised seed lexicons are still necessary.

Finally, when our density-based metric CSLS-D is employed for retrieval, it could achieve further gains in accuracy for most language pairs compared to its counterpart.

5.4 Morphologically Rich Language Results

Søgaard et al. (2018) found that the commonly

	en-et	et-en	en-fi	fi-en	en-el	el-en	en-hu	hu-en	en-pl	pl-en	en-tr	tr-en
MUSE (U+R)	1.7	0.	10.1	59.8	39.1	59.0	50.2	0.1	53.9	0.0	45.4	0.0
5k+Procrustes (R)	31.9	45.6	47.3	59.5	44.6	58.5	53.3	64.8	58.2	66.9	46.3	59.2
id+Procrustes (R)	29.7	40.6	45.0	59.1	40.7	55.1	52.6	63.7	57.3	66.7	45.4	61.4
id+Procrustes (R)*	31.5	-	28.0	-	42.9	-	46.6	-	52.6	-	39.2	-
CSLS (w/o R)	32.9	45.3	47.5	58.7	43.6	57.8	55.3	64.5	59.9	69.1	50.3	60.8
CSLS-D (w/o R)	35.9	45.2	49.4	58.5	45.0	58.1	57.6	64.7	60.9	68.3	52.8	61.6
CSLS (w/ R)	34.4	47.8	48.9	60.2	44.5	61.1	55.1	63.9	59.7	69.0	50.3	60.8
CSLS-D (w/ R)	37.0	47.9	50.7	60.8	46.3	61.6	58.3	64.5	61.5	69.1	53.4	61.2

Table 2: BLI Precision (@1) for morphologically complex languages. **id+Procrustes (R)*** is the result reported in (Søgaard et al., 2018). **5k+Procrustes (R)** uses the training dictionary with 5k unique query words. ('et' is Estonian, 'fi' is Finnish, 'el' is Greek, 'hu' is Hungarian, 'pl' is Persian, 'tr' is Turkish.)

benchmarked languages are morphologically poor isolating or exclusively concatenating languages. They select several languages with different morphological traits and complexity then studied the impacts of language similarities on the bilingual lexicon induction.

They show that a simple trick, harvesting the identical word strings in both languages as an initial dictionary and running the iterative Procrustes analysis described in Sec. 4.2, enables more robust and competitive bilingual induction results over the GAN-based unsupervised method with refinement. We denote this method 'id+Procrustes (R)'.

Tab. 2 shows results on the morphologically complex languages used by Søgaard et al. (2018). For each language pair we run experiments in both directions. The baseline methods in Søgaard et al. (2018) include id+Procrustes (R) and the MUSE (U+R). We run id+Procrustes (R) ourselves and obtain different results from them: except for en-et and en-el, we obtain significantly better results on other language pairs. In addition, we add another strong supervised baseline (5k+Procrustes (R)) with the training dictionary in the MUSE dataset and iterative Procrustes refinement. From Tab. 2, we observe that even without refinement, DeMa-BME (CSLS (w/o R)) outperforms both the unsupervised and supervised baselines on even these difficult morphologically rich languages.

5.5 Cross-lingual Word Similarity

We evaluate DeMa-BWE on the cross-lingual word similarity task from SemEval 2017 (Camacho-Collados et al., 2016) and compare with some strong baselines in Xu et al. (2018). In Tab. 5, following the convention in benchmark evaluation for this task, we report the Pearson correlation scores ($\times 100$). DeMa-BME achieves

Supervised	de-en	es-en	fa-en	it-en
Xing et al. (2015)	72	71	69	72
Shigeto et al. (2015)	72	72	69	72
Artetxe et al. (2016)	73	72	70	73
Artetxe et al. (2017)	70	70	67	71
Unsupervised				
Conneau et al. (2017)	71	71	68	71
Xu et al. (2018)	71	71	67	71
DeMa-BME (w/o R)	72.2	72.2	68.6	72.2

Table 3: Pearson rank correlation ($\times 100$) on cross-lingual word similarity task. Bold indicates the best unsupervised numbers.

the best performance among all the listed unsupervised methods. Compared with the supervised methods, DeMa-BME is also very competitive.

	en-fr	fr-en	en-ja	ja-en
w/o L_{xy} & L_{yx}	45.4	73.3	6.3	27.3
w/o L_{bt}	82.3	82.3	46.2	31.1
w/o L_{sup}	0.1	0.1	0.0	0.0
$\pi(y_j x_i) := \pi(y_j)$	82.1	82.1	46.0	32.5
Full Model	82.6	82.4	46.7	32.9

Table 4: Ablation study on different components of DeMa-BME.

5.6 Ablation Study

Finally, we perform ablation studies in Tab. 4 to examine the effects of different components of DeMa-BWE. In comparison to the full model, we remove the density matching loss L_{xy} & L_{yx} , the weakly supervised loss L_{sup} , the back-translation loss L_{bt} respectively. First, we observe that without the identical strings as the supervised loss, DeMa-BWE fails to converge as the density

matching is difficult given a high-dimensional embedding space to search. Second, when we remove the proposed density matching loss, the model is able to produce reasonable accuracy for fr-en and ja-en, but has undesirable results on en-fr and en-ja, which verifies the necessity of the unsupervised density matching. Third, the back-translation loss is not a crucial component in DeMa-BME; removing it only degrades the model’s performance by a small margin. This indicates that orthogonality is not must-have constraint given the model has enough capacity to learn a good transformation.

In addition, we also compare the frequency-matching based Gaussian mixture weights in (9) with the naive target frequency based weights. As shown in the fourth row of Tab. 4, the performance of DeMa-BWE with the naive weights is nominally worse than the model using the frequency-matching based mixture weights.

6 Conclusion

In this work, we propose a density matching based unsupervised method for learning bilingual word embedding mappings. DeMa-BWE performs well in the task of bilingual lexicon induction. In the future work, we will integrate Gaussian embeddings (Vilnis and McCallum, 2015) with our approach.

Acknowledgments

This work is sponsored by Defense Advanced Research Projects Agency Information Innovation Office (I2O), Program: Low Resource Languages for Emergent Incidents (LORELEI), issued by DARPA/I2O under Contract No. HR0011-15-C-0114. The authors would also like to thank Ruochen Xu, Barun Patra, Joel Ruben Antony Moni for their helpful discussions during drafting this paper.

References

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2016. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2289–2294.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 451–462.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018a. Generalizing and improving bilingual word embedding mappings with a multi-step framework of linear transformations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence (AAAI-18)*.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018b. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *The 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, Melbourne, Australia.

Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018c. Unsupervised statistical machine translation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.

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

Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *arXiv preprint arXiv:1710.04087*.

Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. 2016. Density estimation using real NVP. *arXiv preprint arXiv:1605.08803*.

Manaal Faruqui and Chris Dyer. 2014. Improving vector space word representations using multilingual correlation. In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 462–471, Gothenburg, Sweden. Association for Computational Linguistics.

Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. Bilbowa: Fast bilingual distributed representations without word alignments. In *International Conference on Machine Learning*, pages 748–756.

Edouard Grave, Armand Joulin, and Quentin Berthet. 2018. Unsupervised alignment of embeddings with Wasserstein Procrustes. *arXiv preprint arXiv:1805.11222*.

Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor OK Li. 2018. Universal neural machine translation for extremely low resource languages. In *Meeting of the North American Chapter of the Association for Computational Linguistics (NAACL)*, New Orleans, USA.

- Jiang Guo, Wanxiang Che, David Yarowsky, Haifeng Wang, and Ting Liu. 2015. Cross-lingual dependency parsing based on distributed representations. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, volume 1, pages 1234–1244.
- Aria Haghighi, Percy Liang, Taylor Berg-Kirkpatrick, and Dan Klein. 2008. Learning bilingual lexicons from monolingual corpora. In *Proceedings of ACL-08: HLT*, pages 771–779, Columbus, Ohio. Association for Computational Linguistics.
- Junxian He, Graham Neubig, and Taylor Berg-Kirkpatrick. 2018. Unsupervised learning of syntactic structure with invertible neural projections. In *Proceedings of EMNLP*.
- Karl Moritz Hermann and Phil Blunsom. 2014. Multilingual models for compositional distributed semantics. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 58–68.
- Yedid Hoshen and Lior Wolf. 2018. Non-adversarial unsupervised word translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 469–478.
- Pratik Jawanpuria, Arjun Balgovind, Anoop Kunchukuttan, and Bamdev Mishra. 2018. Learning multilingual word embeddings in latent metric space: a geometric approach. *arXiv preprint arXiv:1808.08773*.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. Loss in translation: Learning bilingual word mapping with a retrieval criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2984.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Diederik P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. 2016. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems*, pages 4743–4751.
- Durk P Kingma and Prafulla Dhariwal. 2018. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10235–10244.
- Alexandre Klementiev, Ivan Titov, and Binod Bhattarai. 2012. Inducing crosslingual distributed representations of words. *Proceedings of COLING 2012*, pages 1459–1474.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- George Papamakarios, Iain Murray, and Theo Pavlakou. 2017. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347.
- Miloš Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2010. Hubs in space: Popular nearest neighbors in high-dimensional data. *J. Mach. Learn. Res.*, 11:2487–2531.
- Danilo Jimenez Rezende and Shakir Mohamed. 2015. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning—Volume 37*, pages 1530–1538. JMLR. org.
- Sebastian Ruder, Ivan Vulic, and Anders Søgaard. 2017. A survey of cross-lingual embedding models. *CoRR, abs/1706.04902*.
- Yutaro Shigeto, Ikumi Suzuki, Kazuo Hara, Masashi Shimbo, and Yuji Matsumoto. 2015. Ridge regression, hubness, and zero-shot learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 135–151. Springer.
- Samuel L Smith, David HP Turban, Steven Hamblin, and Nils Y Hammerla. 2017. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *arXiv preprint arXiv:1702.03859*.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. On the limitations of unsupervised bilingual dictionary induction. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788. Association for Computational Linguistics.
- Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual models of word embeddings: An empirical comparison. *arXiv preprint arXiv:1604.00425*.
- Luke Vilnis and Andrew McCallum. 2015. Word representations via gaussian embedding. *International Conference on Learning Representations*.
- Laura Wendlandt, Jonathan K. Kummerfeld, and Rada Mihalcea. 2018. Factors influencing the surprising instability of word embeddings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2092–2102. Association for Computational Linguistics.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 2015. Normalized word embedding and orthogonal transform for bilingual word translation. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1006–1011.

- Ruo Chen Xu, Yiming Yang, Naoki Otani, and Yuexin Wu. 2018. Unsupervised cross-lingual transfer of word embedding spaces. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Brussels, Belgium.
- Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. Earth mover’s distance minimization for unsupervised bilingual lexicon induction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1934–1945.
- Zhisong Zhang, Wasil Uddin Ahmad, Xuezhe Ma, Eduard Hovy, Kai-Wei Chang, and Nanyun Peng. 2018. Near or far, wide range zero-shot cross-lingual dependency parsing. *arXiv preprint arXiv:1811.00570*.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. Transfer learning for low-resource neural machine translation. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*.