

Learning Translations via Matrix Completion

Derry Wijaya¹, Brendan Callahan¹, John Hewitt¹, Jie Gao¹, Xiao Ling¹,
Marianna Apidianaki^{1,2} and Chris Callison-Burch¹

¹Computer and Information Science Department, University of Pennsylvania

²LIMSI, CNRS, Université Paris-Saclay, 91403 Orsay

derry@seas.upenn.edu

Abstract

Bilingual Lexicon Induction is the task of learning word translations without bilingual parallel corpora. We model this task as a matrix completion problem, and present an effective and extendable framework for completing the matrix. This method harnesses diverse bilingual and monolingual signals, each of which may be incomplete or noisy. Our model achieves state-of-the-art performance for both high and low resource languages.

1 Introduction

Machine translation (MT) models typically require large, sentence-aligned bilingual texts to learn good translation models (Wu et al., 2016; Sennrich et al., 2016a; Koehn et al., 2003). However, for many language pairs, such parallel texts may only be available in limited quantities, which is problematic. Alignments at the word- or subword- levels (Sennrich et al., 2016b) can be inaccurate in the limited parallel texts, which can in turn lead to inaccurate translations. Due to the low quantity and thus coverage of the texts, there may still be “out-of-vocabulary” words encountered at run-time. The Bilingual Lexicon Induction (BLI) task (Rapp, 1995), which learns word translations from monolingual or comparable corpora, is an attempt to alleviate this problem. The goal is to use plentiful, more easily obtainable, monolingual or comparable data to infer word translations and reduce the need for parallel data to learn good translation models. The word translations obtained by BLI can, for example, be used to augment MT systems and improve alignment accuracy, coverage, and translation quality (Gulcehre et al., 2016; Callison-Burch et al., 2006; Daumé and Jagarlamudi, 2011).




Figure 1: Our framework allows us to use a diverse range of signals to learn translations, including incomplete bilingual dictionaries, information from related languages (like Indonesian loan words from Dutch shown here), word embeddings, and even visual similarity cues.

Previous research has explored different sources for estimating translation equivalence from monolingual corpora (Schafer and Yarowsky, 2002; Klementiev and Roth, 2006; Irvine and Callison-Burch, 2013, 2017). These monolingual signals, when combined in a supervised model, can enhance end-to-end MT for low resource languages (Klementiev et al., 2012a; Irvine and Callison-Burch, 2016). More recently, similarities between words in different languages have been approximated by constructing a shared bilingual word embedding space with different forms of bilingual supervision (Upadhyay et al., 2016).

We present a framework for learning translations by combining diverse signals of translation that are each potentially sparse or noisy. We use matrix factorization (MF), which has been shown to be effective for harnessing incomplete or noisy distant supervision from multiple sources of information (Fan et al., 2014; Rocktäschel et al., 2015). MF is also shown to result in good cross-lingual representations for tasks such as alignment (Goutte et al., 2004), QA (Zhou et al., 2013), and cross-lingual word embeddings (Shi et al., 2015).

Specifically, we represent translation as a matrix with source words in the columns and target words in the rows, and model the task of learning translations as a matrix completion problem. Starting from some observed translations (e.g., from existing bilingual dictionaries,) we infer missing translations in the matrix using MF with a Bayesian Personalized Ranking (BPR) objective (Rendle et al., 2009). We select BPR for a number of reasons: (1) BPR has been shown to outperform traditional supervised methods in the presence of positive-only data (Riedel et al., 2013), which is true in our case since we only observe positive translations. (2) BPR is easily extendable to incorporate additional signals for inferring missing values in the matrix (He and McAuley, 2016). Since observed translations may be sparse, i.e. the “cold start” problem in the matrix completion task, incorporating additional signals of translation equivalence estimated on monolingual corpora is useful. (3) BPR is also shown to be effective for multilingual transfer learning (Verga et al., 2016). For low resource source languages, there may be related, higher resource languages from which we can project available translations (e.g., translations of loan words) to the target language (Figure 1).

We conduct large scale experiments to learn translations from both low and high resource languages to English and achieve state-of-the-art performance on these languages. Our main contributions are as follows:

- We introduce a MF framework that learns translations by integrating diverse bilingual and monolingual signals of translation, each potentially noisy/incomplete.
- The framework is easily extendable to incorporate additional signals of translation equivalence. Since ours is a framework for integration, each signal can be improved separately to improve the overall system.
- Large scale experiments on both low and high resource languages show the effectiveness of our model, outperforming the current state-of-the-art.
- We make our code, datasets, and output translations publicly available.¹

2 Related Work

Bilingual Lexicon Induction Previous research has used different sources for estimating transla-

tions from monolingual corpora. Signals such as contextual, temporal, topical, and orthographic similarities between words are used to measure their translation equivalence (Schafer and Yarowsky, 2002; Klementiev and Roth, 2006; Irvine and Callison-Burch, 2013, 2017).

With the increasing popularity of word embeddings, many recent works approximate similarities between words in different languages by constructing a shared bilingual embedding space (Klementiev et al., 2012b; Zou et al., 2013; Vulić and Moens, 2013; Mikolov et al., 2013a; Faruqui and Dyer, 2014; Chandar A P et al., 2014; Gouws et al., 2015; Luong et al., 2015; Lu et al., 2015; Upadhyay et al., 2016). In the shared space, words from different languages are represented in a language-independent manner such that similar words, regardless of language, have similar representations. Similarities between words can then be measured in the shared space. One approach to induce this shared space is to learn a mapping function between the languages’ monolingual semantic spaces (Mikolov et al., 2013a; Dinu et al., 2014). The mapping relies on seed translations which can be from existing dictionaries or be reliably chosen from pseudo-bilingual corpora of comparable texts e.g., Wikipedia with interlanguage links. Vulić and Moens (2015) show that by learning a linear function with a reliably chosen seed lexicon, they outperform other models with more expensive bilingual signals for training on benchmark data.

Most prior work on BLI however, either makes use of only one monolingual signal or uses unsupervised methods (e.g., rank combination) to aggregate the signals. Irvine and Callison-Burch (2016) show that combining monolingual signals in a supervised logistic regression model produces higher accuracy word translations than unsupervised models. More recently, Vulić et al. (2016) show that their multi-modal model that employs a simple weighted-sum of word embeddings and visual similarities can improve translation accuracy. These works show that there is a need for combining diverse, multi-modal monolingual signals of translations. In this paper, we take this step further by combining the monolingual signals with bilingual signals of translations from existing bilingual dictionaries of related, “third” languages.

Bayesian Personalized Ranking (BPR) Our approach is based on extensions to the probabilis-

¹<http://www.cis.upenn.edu/%7Ederry/translations.html>

tic model of MF in collaborative filtering (Koren et al., 2009; Rendle et al., 2009). We represent our translation task as a matrix with source words in the columns and target words in the rows (Figure 1). Based on some observed translations in the matrix found in a seed dictionary, our model learns low-dimensional feature vectors that encode the latent properties of the words in the row and the words in the column. The dot product of these vectors, which indicate how “aligned” the source and the target word properties are, captures how likely they are to be translations.

Since we do not observe false translations in the seed dictionary, the training data in the matrix consists only of positive translations. The absence of values in the matrix does not imply that the corresponding words are not translations. In fact, we seek to predict which of these missing values are true. The BPR approach to MF (Rendle et al., 2009) formulates the task of predicting missing values as a ranking task. With the assumption that observed true translations should be given higher values than unobserved translations, BPR learns to optimize the difference between values assigned to the observed translations and values assigned to the unobserved translations.

However, due to the sparsity of existing bilingual dictionaries (for some language pairs such dictionaries may not exist), the traditional formulation of MF with BPR suffers from the “cold start” issue (Gantner et al., 2010; He and McAuley, 2016; Verga et al., 2016). In our case, these are situations in which some source words have no translations to any word in the target or related languages. For these words, additional information, e.g., monolingual signals of translation equivalence or language-independent representations such as visual representations, must be used.

We use bilingual translations from the source to the target language, English, obtained from Wikipedia page titles with interlanguage links. Since Wikipedia pages in the source language may be linked to pages in languages other than English, we also use high accuracy, crowdsourced translations (Pavlick et al., 2014) from these *third* languages to English as additional bilingual translations. To alleviate the cold start issue, when a source word has no existing known translation to English or other third languages, our model backs-off to additional signals of translation equivalence estimated based on its word embedding and visual

representations.

3 Method

In this section, we describe our framework for integrating bilingual and monolingual signals for learning translations. First we formulate the task of Bilingual Lexicon Induction, and introduce our model for learning translations given observed translations and additional monolingual/language-independent signals. Then we derive our learning procedure using the BPR objective function.

Problem Formulation Given a set of source words F , a set of target words E , the pair $\langle e, f \rangle$ where $e \in E$ and $f \in F$ is a candidate translation with an associated score $x_{e,f} \in [0, 1]$ indicating the confidence of the translation. The input to our model is a set of observed translations $T := \{\langle e, f \rangle \mid x_{e,f} = 1\}$. These could come from an incomplete bilingual dictionary. We also add word identities to the matrix i.e., we define $T^{identity} := \{\langle e, e \rangle\}$, where $T^{identity} \subset T$. The task of Bilingual Lexicon Induction is then to generate *missing* translations: for a given source word f and a set of target words $\{e \mid \langle e, f \rangle \notin T\}$, predict the score $x_{e,f}$ of how likely it is for e to be a translation of f .

Bilingual Signals for Translation One way to predict $x_{e,f}$ is by using matrix factorization. The problem of predicting $x_{e,f}$ can be seen as a task of estimating a matrix $X : E \times F$. X is approximated by a matrix product of two low-rank matrices $P : |E| \times k$ and $Q : |F| \times k$:

$$\hat{X} := PQ^T$$

where k is the rank of the approximation. Each row p_e in P can be seen as a feature vector describing the latent properties of the target word e , and each row q_f of Q describes the latent properties of the source word f . Their dot product encodes how aligned the latent properties are and, since these vectors are trained on observed translations, it encodes how likely they are to be translation of each other. Thus, we can write this formulation of predicted scores $\hat{x}_{e,f}$ with MF as:

$$\hat{x}_{e,f}^{MF} = p_e^T q_f = \sum_{i=1}^k p_{ei} \cdot q_{fi} \quad (1)$$

Auxiliary Signals for Translation Because the observed bilingual translations may be sparse, the

										θ_f
	○	○	○	•	•	•	•	•	•	
	○	○	○	•	•	•	•	•	•	
	
	○	○	○	•	•	•	•	•	•	
	○	○	○	•	•	•	•	•	•	
		fridge	house	sleep	kulkas	koki	tidur			
EN	fridge	■								
	house		■							
	sleep			■						
	hijzelf									
NL	koelkast	■			■					
	kokkien					■				

Figure 2: The word *tidur* (id) is a cold word with no associated translation in the matrix. Auxiliary features θ_f about the words can be used to predict translations for cold words.

MF approach can suffer from the existence of *cold* items: words that have none or too few associated observed translations to estimate their latent dimensions accurately (Figure 2). Additional signals for measuring translation equivalence can alleviate this problem. Hence, in the case of cold words, we use a formulation of $\hat{x}_{e,f}$ that involves auxiliary features about the words in the predicted $\hat{x}_{e,f}$:

$$\hat{x}_{e,f}^{\text{AUX}} = \theta_e^T \theta_f + \beta^T \theta_f \quad (2)$$

θ_f represents an auxiliary information about the cold word f e.g., its word embedding or visual features. θ_e is a feature vector to be trained, whose dot product with θ_f models the extent to which the word e matches the auxiliary features of word f . In practice, learning θ_e amounts to learning a classifier, one for each target word e that learns weights θ_e given the feature vectors θ_f of its translations. β models the targets' overall bias toward a given word f .

Since each word can have multiple additional feature vectors, we can formulate $\hat{x}_{e,f}^{\text{AUX}}$ as a weighted sum of available auxiliary features²:

$$\hat{x}_{e,f}^{\text{AUX}} = \alpha_1 \theta_e^T \theta_f + \alpha_2 \gamma_e^T \gamma_f + \dots + \alpha_n \delta_e^T \delta_f$$

where α_m are parameters assigned to control the contribution of each auxiliary feature.

In practice, we can combine the MF and auxiliary formulations by defining:

$$\hat{x}_{e,f} = \hat{x}_{e,f}^{\text{MF}} + \hat{x}_{e,f}^{\text{AUX}}$$

²We omit bias terms for brevity.

However, since bilingual signals that are input to $\hat{x}_{e,f}^{\text{MF}}$ are often precise but sparse, while monolingual signals that are input to $\hat{x}_{e,f}^{\text{AUX}}$ are often noisy and not sparse, in our model we only back-off to the less precise $\hat{x}_{e,f}^{\text{AUX}}$ for cold source words that have none or too few associated translations (more details are given in the experiments, Section 4). For other source words, we use $\hat{x}_{e,f}^{\text{MF}}$ to predict.

Learning with Bayesian Personalized Ranking

Unlike traditional supervised models that try to maximize the scores assigned to positive instances (in our case, observed translations), the objective of Bayesian Personalized Ranking (BPR) is to maximize the *difference* in scores assigned to the observed translations compared to those assigned to the unobserved translations. Given a training set D consisting of triples of the form $\langle e, f, g \rangle$, where $\langle e, f \rangle \in T$ and $\langle e, g \rangle \notin T$, BPR wants to maximize $\hat{x}_{e,f,g}$, defined as:

$$\hat{x}_{e,f,g} = \hat{x}_{e,f} - \hat{x}_{e,g}$$

where $\hat{x}_{e,f}$ and $\hat{x}_{e,g}$ can be defined either by eq. 1 or eq. 2 (for cold words). Specifically, BPR optimizes (Rendle et al., 2009):

$$\sum_{\langle e,f,g \rangle \in D} \ln \sigma(\hat{x}_{e,f,g}) - \lambda_{\Theta} \|\Theta\|^2$$

where σ is the logistic sigmoid function, Θ is the parameter vector of $\hat{x}_{e,f,g}$ to be trained, and λ_{Θ} is its hyperparameter vector. BPR can be trained using stochastic gradient ascent where a triple $\langle e, f, g \rangle$ is sampled from D and parameter updates are performed:

$$\Theta \leftarrow \Theta + \eta \cdot (\sigma(-\hat{x}_{e,f,g}) \frac{\partial \hat{x}_{e,f,g}}{\partial \Theta} - \lambda_{\Theta} \Theta)$$

η is the learning rate. Hence, for the MF formulation of $\hat{x}_{e,f,g}$, we can sample a triple $\langle e, f, g \rangle$ from D and update its parameters as:

$$p_e \leftarrow p_e + \eta \cdot (\sigma(-\hat{x}_{e,f,g}^{\text{MF}})(q_f - q_g) - \lambda_P p_e)$$

$$q_f \leftarrow q_f + \eta \cdot (\sigma(-\hat{x}_{e,f,g}^{\text{MF}})(p_e) - \lambda_{Q^+} q_f)$$

$$q_g \leftarrow q_g + \eta \cdot (\sigma(-\hat{x}_{e,f,g}^{\text{MF}})(-p_e) - \lambda_{Q^-} q_g)$$

while for the auxiliary formulation of $\hat{x}_{e,f,g}$, we can sample a triple $\langle e, f, g \rangle$ from D and update its parameters as:

$$\theta_e \leftarrow \theta_e + \eta \cdot (\sigma(-\hat{x}_{e,f,g}^{\text{AUX}})(\theta_f - \theta_g) - \lambda_{\Theta} \theta_e)$$

$$\beta \leftarrow \beta + \eta \cdot (\sigma(-\hat{x}_{e,f,g}^{\text{AUX}})(\theta_f - \theta_g) - \lambda_{\beta} \beta)$$

4 Experiments

To implement our approach, we extend the implementation of BPR in LIBREC³ which is a publicly available Java library for recommender systems.

We evaluate our model for the task of Bilingual Lexicon Induction (BLI). Given a source word f , the task is to rank all candidate target words e by their predicted translation scores $\hat{x}_{e,f}$. We conduct large-scale experiments on 27 low- and high-resource source languages and evaluate their translations to English. We use the 100K most frequent words from English Wikipedia as candidate English target words (E).

At test time, for each source language, we evaluate the top-10 accuracy (Acc_{10}): the percent of source language words in the test set for which a correct English translation appears in the top-10 ranked English candidates.

4.1 Data

4.1.1 Test sets

We use benchmark test sets for the task of bilingual lexicon induction to evaluate the performance of our model. The **VULIC1000** dataset (Vulic and Moens, 2016) comprises 1000 nouns in Spanish, Italian, and Dutch, along with their one-to-one ground-truth word translations in English.

We construct a new test set (**CROWDTEST**) for a larger set of 27 languages from crowdsourced dictionaries (Pavlick et al., 2014). For each language, we randomly pick up to 1000 words that have only one English word translation in the crowdsourced dictionary to be the test set for that language. On average, there are 967 test source words with a variety of POS per language. Since different language treats grammatical categories such as tense and number differently (for example, unlike English, tenses are not expressed by specific forms of words in Indonesian (id); rather, they are expressed through context), we make our evaluation on all languages in **CROWDTEST** generic by treating a predicted English translation of a foreign word as correct as long as it has the same lemma as the gold English translation. To facilitate further research, we make **CROWDTEST** publicly available in our website.

4.1.2 Bilingual Signals for Translation

We use Wikipedia to incorporate information from a *third* language into the matrix, with ob-




Figure 3: Wikipedia pages with observed translations to the source (id) and the target (en) languages act as a *third* language in the matrix.

served translations to both the source language and the target language, English. We first collect all interlingual links from English Wikipedia pages to pages in other languages. Using these links, we obtain translations of Wikipedia page titles in many languages to English e.g., `id.wikipedia.org/wiki/Kulkas` \rightarrow `fridge` (en). The observed translations are projected to fill the missing translations in the matrix (Figure 3). We call these bilingual translations **WIKI**.

From the links that we have collected, we can also infer links from Wikipedia pages in the source language to other pages in non-target languages e.g., `id.wikipedia.org/wiki/Kulkas` \rightarrow `it.wikipedia.org/wiki/Frigorifero`. The titles of these pages can be translated to English if they exist as entries in the dictionaries. These non-source, non-target language pages can act as yet another third language whose observed translations can be projected to fill the missing translations in the matrix (Figure 3). We call these bilingual translations **WIKI+CROWD**.

4.1.3 Monolingual Signals for Translation

We define *cold* source words in our experiments as source words that have no associated **WIKI** translations and fewer than 2 associated **WIKI+CROWD** translations. For each cold source word f , we predict the score of its translation to each candidate English word e using the auxiliary formulation of $\hat{x}_{e,f}$ (Equation 2). There are two auxiliary signals about the words that we use in our experiments: (1) bilingually informed word embeddings and (2) visual representations.

Bilingually Informed Word Embeddings For each language, we learn monolingual embeddings for its words by training a standard monolingual `word2vec` skipgram model (Mikolov

³<https://www.librec.net/index.html>

et al., 2013b) on tokenized Wikipedia pages of that language using Gensim (Řehůřek and Sojka, 2010). We obtain 100-dimensional word embeddings with 15 epochs, 15 negatives, window size of 5, and cutoff value of 5.

Given two monolingual embedding spaces \mathbb{R}^{d_F} and \mathbb{R}^{d_E} of the source and target languages F and E , where d_f and d_e denote the dimensionality of the monolingual embedding spaces, we use the set of crowdsourced translations that are not in the test set as our seed bilingual translations⁴ and learn a mapping function $\mathbf{W} \in \mathbb{R}^{d_E \times d_F}$ that maps the target language vectors in the seed translations to their corresponding source language vectors.⁵

We learn two types of mapping: linear and non-linear, and compare their performances. The linear mapping (Mikolov et al., 2013a; Dinu et al., 2014) minimizes: $\|\mathbf{X}_E \mathbf{W} - \mathbf{X}_F\|_F^2$ where, following the notation in (Vulić and Korhonen, 2016), \mathbf{X}_E and \mathbf{X}_F are matrices obtained by respective concatenation of target language and source language vectors that are in the seed bilingual translations. We solve this optimization problem using stochastic gradient descent (SGD).

We also consider a non-linear mapping (Socher et al., 2013) using a simple four-layer neural network, $\mathbf{W} = (\phi^{(1)}, \phi^{(2)}, \phi^{(3)}, \phi^{(4)})$ that is trained to minimize:

$$\sum_{\mathbf{x}_f \in \mathbf{X}_F} \sum_{\mathbf{x}_e \in \mathbf{X}_E} \|\mathbf{x}_f - \phi^{(4)}(s(\phi^{(3)}(s(\phi^{(2)}(s(\phi^{(1)}\mathbf{x}_e))))))\|^2$$

where $\phi^{(1)} \in \mathbb{R}^{h_1 \times d_E}$, $\phi^{(2)} \in \mathbb{R}^{h_2 \times h_1}$, $\phi^{(3)} \in \mathbb{R}^{h_3 \times h_2}$, $\phi^{(4)} \in \mathbb{R}^{d_F \times h_3}$, h_n is the size of the hidden layer, and $s = \tanh$ is the chosen non-linear function.

Once the map \mathbf{W} is learned, all candidate target word vectors \mathbf{x}_e can be mapped into the source language embedding space \mathbb{R}^{d_F} by computing $\mathbf{x}_e^T \mathbf{W}$. Instead of the raw monolingual word embeddings \mathbf{x}_e , we use these bilingually-informed mapped word vectors $\mathbf{x}_e^T \mathbf{W}$ as the auxiliary word features **WORD-AUX** to estimate $\hat{x}_{e,f}^{\text{AUX}}$.

Visual Representations Pilot Study Recent work (Vulić et al., 2016) has shown that combining word embeddings and visual representations of words can help achieve more accurate bilingual translations. Since the visual representation of a



Figure 4: Five images for the French word *eau* and its top 4 translations ranked using visual similarities of images associated with English words (Bergsma and Van Durme, 2011)

word seems to be language-independent (e.g. the concept of *water* has similar images whether expressed in English or French (Figure 4), the visual representations of a word may be useful for inferring its translation and for complementing the information learned in text.

We performed a pilot study to include visual features as auxiliary features in our framework. We use a large multilingual corpus of labeled images (Callahan, 2017) to obtain the visual representation of the words in our source and target languages. The corpus contains 100 images for up to 10k words in each of 100 foreign languages, plus images of each of their translations into English. For each of the images, a convolutional neural network (CNN) feature vector is also provided following the method of Kiela et al. (2015). For each word, we use 10 images provided by this corpus and use their CNN features as auxiliary visual features **VISUAL-AUX** to estimate $\hat{x}_{e,f}^{\text{AUX}}$.

4.1.4 Combining Signals

During training, we trained the parameters of $\hat{x}_{e,f}^{\text{MF}}$ and $\hat{x}_{e,f}^{\text{AUX}}$ using a variety of signals:

- $\hat{x}_{e,f}^{\text{MF-W}}$ is trained using **WIKI** translations as the set of observed translations T
- $\hat{x}_{e,f}^{\text{MF-W+C}}$ is trained using **WIKI+CROWD** translations as the set of observed T
- $\hat{x}_{e,f}^{\text{AUX-WE}}$ is trained using the set of word identities T^{identity} and **WORD-AUX** as θ_f
- $\hat{x}_{e,f}^{\text{AUX-VIS}}$ is trained using the set of word identities T^{identity} and **VISUAL-AUX** as θ_f

During testing, we use the following back-off scheme to predict translation scores given a source word f and a candidate target word e :

⁴On average, there are 9846 crowdsourced translations per language that we can use as seed translations.

⁵We find that mapping from target to source vectors gives better performances across models in our experiments.

$$\hat{x}_{e,f} = \begin{cases} \hat{x}_{e,f}^{\text{MF-W}} & \text{if } f \text{ has } \geq 1 \text{ associated } \mathbf{WIKI}, \\ \hat{x}_{e,f}^{\text{MF-W+C}} & \text{else if } f \text{ has } \geq 2 \\ & \text{associated } \mathbf{WIKI+CROWD}, \\ \hat{x}_{e,f}^{\text{AUX}} & \text{otherwise} \end{cases}$$

where $\hat{x}_{e,f}^{\text{AUX}} = \alpha_{\text{we}} \hat{x}_{e,f}^{\text{AUX-WE}} + \alpha_{\text{vis}} \hat{x}_{e,f}^{\text{AUX-VIS}}$

4.2 Results

We conduct experiments using the following variants of our model, each of which progressively incorporates more signals to rank candidate English target words. When a variant uses more than one formulation of $\hat{x}_{e,f}$, it applies them using the back-off scheme that we have described before.

- **BPR_W** uses *only* $\hat{x}_{e,f}^{\text{MF-W}}$
- **BPR_W+C** uses $\hat{x}_{e,f}^{\text{MF-W}}$ and $\hat{x}_{e,f}^{\text{MF-W+C}}$
- **BPR_LN** uses *only* $\hat{x}_{e,f}^{\text{AUX-WE}}$ with linear mapping
- **BPR_NN** uses *only* $\hat{x}_{e,f}^{\text{AUX-WE}}$ with neural network (NN) mapping
- **BPR_WE** uses $\hat{x}_{e,f}^{\text{MF-W}}$, $\hat{x}_{e,f}^{\text{MF-W+C}}$, and $\hat{x}_{e,f}^{\text{AUX-WE}}$ with NN mapping
- **BPR_VIS** adds $\hat{x}_{e,f}^{\text{AUX-VIS}}$ to **BPR_WE**

Table 1: Acc_{10} performance on **VULIC1000**

	Baseline (MNN)	BPR+MNN	BPR_LN	BPR_WE
IT-EN	78.8%	79.4%	81.3%	86.0%
ES-EN	81.8%	82.1%	83.4%	87.1%
NL-EN	80.8%	81.6%	83.2%	87.2%

We evaluate the performance of **BPR_WE** against a baseline that is the state-of-the-art model of Vulic and Korhonen (2016), on benchmark **VULIC1000** (Table 1). The baseline (**MNN**) learns a linear mapping between monolingual embedding spaces and finds translations in an *unsupervised* manner: it ranks candidate target words based on their cosine similarities to the source word in the mapped space. As seed translation pairs, **MNN** uses mutual nearest neighbor pairs (MNN) obtained from *pseudo-bilingual* corpora constructed from unannotated monolingual data of the source and target languages (Vulic and Moens, 2016). We train **MNN** and our models using the same 100-dimensional `word2vec` monolingual word embeddings.

As seen in Table 1, we see the benefit of learning translations in a supervised manner.

BPR+MNN uses the same MNN seed translations as **MNN**, obtained from unannotated monolingual data of English and the foreign language, to learn the linear mapping between their embedding spaces. However, unlike **MNN**, **BPR+MNN** uses the mapped word vectors to predict ranking in a *supervised* manner with BPR objective. This results in higher accuracies than **MNN**. Using seed translations from crowdsourced dictionaries to learn the linear mapping (**BPR_LN**) improves accuracies even further compared to using MNN seed translations obtained from unannotated data. Finally, **BPR_WE** that learns translations in a supervised manner and uses third language translations and non-linear mapping (trained with crowdsourced translations not in the test set) performs consistently and very significantly better than the state-of-the-art on all benchmark test sets. This shows that incorporating more and better signals of translation can improve performance significantly.

Evaluating on **CROWDTEST**, we observe a similar trend over all 27 languages (Figure 5). Particularly, we see that **BPR_W** and **BPR_W+C** suffer from the *cold start* issue where there are too few or no observed translations in the matrix to make accurate predictions. Incorporating auxiliary information in the form of bilingually-informed word embeddings improves the accuracy of the predictions dramatically. For many languages, learning these bilingually-informed word embeddings with non-linear mapping improves accuracy even more. The top accuracy scores achieved by the model vary across languages and seem to be influenced by the amount of data i.e., Wikipedia tokens and seed lexicons entries available for training. Somali (so) for example, has only 0.9 million tokens available in its Wikipedia for training the `word2vec` embeddings and only 3 thousand seed translations for learning the mapping between the word embedding spaces. In comparison, Spanish (es) has over 500 million tokens available in its Wikipedia and 11 thousand seed translations. We also believe that our choice of tokenization may not be suitable for some languages – we use a simple regular-expression based tokenizer for many languages that do not have a trained NLTK⁶ tokenization model. This may influence performance on languages such as Vietnamese (vi) on which we have a low performance despite its large Wikipedia corpus.

⁶<http://www.nltk.org/>




Figure 5: Acc_{10} on **CROWDTEST** across all 27 languages show that adding more and better signals for translation improves translation accuracies. The top accuracies achieved by our model: **BPR_WE** vary across languages and appear to be influenced by the amount of data (Wikipedia tokens and seed translations) and tokenization available for the language.

Table 2: Top-5 translations of the Indonesian word *kesadaran* (*awareness*) using different model variants

BPR_W	BPR_LN	BPR_NN	BPR_WE
<i>kesadaran</i>	<i>kesadaran</i>	<i>kesadaran</i>	<i>kesadaran</i>
consciousness	consciousness	consciousness	conscience
goddess	<i>awareness</i>	empathy	<i>awareness</i>
friendship	empathy	<i>awareness</i>	understanding
night	perception	perceptions	consciousness
nation	mindedness	perception	acquaintance

Some example translations of an Indonesian word produced by different variants of our model are shown in Table 2. Adding third language translation signals on top of the bilingually-informed auxiliary signals improves accuracies even further.⁷ The accuracies achieved by **BPR_WE** on these languages are significantly better than previously reported accuracies (Irvine and Callison-Burch, 2017) on test sets constructed from the same crowdsourced dictionaries (Pavlick et al., 2014)⁸.

The accuracies across languages appear to improve consistently with the amount of signals being input to the model. In the following experiments, we investigate how sensitive these improvements are with varying training size.

In Figure 6, we show accuracies obtained by

⁷Actual improvement per language depends on the coverage of the Wikipedia interlanguage links for that language

⁸The comparison however, cannot be made apples-to-apples since the way Irvine and Callison-Burch (2017) select test sets from the crowdsourced dictionaries maybe different and they do not release the test sets

BPR_WE with varying sizes of seed translation lexicons used to train its mapping. The results show that a seed lexicon size of 5K is enough across languages to achieve optimum performance. This finding is consistent with the finding of Vulić and Korhonen (2016) that accuracies peak at about 5K seed translations across all their models and languages. For future work, it will be interesting to investigate further why this is the case: e.g., how optimal seed size is related to the quality of the seed translations and the size of the test set, and how the optimum seed size should be chosen. Lastly, we experiment with incorporating auxiliary visual signals for learning translations on the multilingual image corpus (Callahan, 2017). The corpus contains 100 images for up to 10K words in each of 100 foreign languages, plus images of each of their translations into English. We train and test our **BPR_VIS** model to learn translations of 5 low- and high-resource languages in this corpus. We use the translations of up to 10K words in each of these languages as test set and use up to




Figure 6: Acc_{10} across different seed lexicon sizes

Table 3: Acc_{10} performance on the multilingual image corpus test set (Callahan, 2017)

	Baseline (CNN-AvgMax)	BPR_VIS	# Seeds
IT-EN	31.4%	55.8%	581
ES-EN	33.0%	58.3%	488
NL-EN	35.5%	69.2%	1857
FR-EN	37.1%	65.9%	1697
ID-EN	36.9%	45.3%	462

10 images (CNN features) of the words in this set as auxiliary visual signals to predict their translations. In this experiment, we weigh auxiliary word embedding and visual features equally. To train the mapping of our word embedding features, we use as seeds crowdsourced translations not in test set.

We compare the quality of our translations with the baseline CNN-AVGMAX (Bergsma and Van Durme, 2011), which considers cosine similarities between individual images from the source and target word languages and takes average of their maximum similarities as the final similarity between a source and a target word. For each source word, the candidate target words are ranked according to these final similarities. This baseline has been shown to be effective for inducing translations from images, both in the uni-modal (Bergsma and Van Durme, 2011; Kiela et al., 2015) and multi-modal models (Vulić et al., 2016).

As seen in Table 3, incorporating additional bilingual and textual signals to the visual signals improves translations. Accuracies on these image corpus’ test sets are lower overall as they contain a lot of translations from our crowdsourced dictionaries; thus we have much less seeds to train our word embedding mapping. Furthermore, these test sets contain 10 times as many translations as our previous test sets. Using more images instead of just 10 per word may also improve performance.

5 Conclusion

In this paper, we propose a novel framework for combining diverse, sparse and potentially noisy multi-modal signals for translations. We view the problem of learning translations as a matrix completion task and use an effective and extendable matrix factorization approach with BPR to learn translations.

We show the effectiveness of our approach in large scale experiments. Starting from minimally-trained monolingual word embeddings, we consistently and very significantly outperform state-of-the-art approaches by combining these features with other features in a supervised manner using BPR. Since our framework is modular, each input to our prediction can be improved separately to improve the whole system e.g., by learning better word embeddings or a better mapping function to input into the auxiliary component. Our framework is also easily extendable to incorporate more bilingual and auxiliary signals of translation equivalence.

Acknowledgments

This material is based in part on research sponsored by DARPA under grant number HR0011-15-C-0115 (the LORELEI program). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes. The views and conclusions contained in this publication are those of the authors and should not be interpreted as representing official policies or endorsements of DARPA and the U.S. Government.

This work was also supported by the French National Research Agency under project ANR-16-CE33-0013, and by Amazon through the Amazon Academic Research Awards (AARA) program.

References

- Shane Bergsma and Benjamin Van Durme. 2011. Learning Bilingual Lexicons Using the Visual Similarity of Labeled Web Images. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, pages 1764–1769, Barcelona, Spain.
- Brendan Callahan. 2017. Image-based bilingual lexicon induction for low resource languages. Master’s thesis, University of Pennsylvania.
- Chris Callison-Burch, Philipp Koehn, and Miles Osborne. 2006. Improved Statistical Machine Translation Using Paraphrases. In *Proceedings of the*

- Human Language Technology Conference of the NAACL, Main Conference*, pages 17–24, New York City.
- Sarath Chandar A P, Stanislas Lauly, Hugo Larochelle, Mitesh Khapra, Balaraman Ravindran, Vikas C Raykar, and Amrita Saha. 2014. An Autoencoder Approach to Learning Bilingual Word Representations. In *Advances in Neural Information Processing Systems 27*, pages 1853–1861.
- Hal Daumé and Jagadeesh Jagarlamudi. 2011. Domain Adaptation for Machine Translation by Mining Unseen Words. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 407–412, Portland, Oregon.
- Georgiana Dinu, Angeliki Lazaridou, and Marco Baroni. 2014. Improving zero-shot learning by mitigating the hubness problem. In *Proceedings of ICLR Workshop*, San Diego, California.
- Miao Fan, Deli Zhao, Qiang Zhou, Zhiyuan Liu, Thomas Fang Zheng, and Edward Y. Chang. 2014. Distant Supervision for Relation Extraction with Matrix Completion. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 839–849, Baltimore, Maryland.
- 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.
- Zeno Gantner, Lucas Drumond, Christoph Freudenthaler, Steffen Rendle, and Lars Schmidt-Thieme. 2010. Learning attribute-to-feature mappings for cold-start recommendations. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 176–185. IEEE.
- Cyril Goutte, Kenji Yamada, and Eric Gaussier. 2004. Aligning words using matrix factorisation. In *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, page 502. Association for Computational Linguistics.
- Stephan Gouws, Yoshua Bengio, and Greg Corrado. 2015. BilBOWA: Fast Bilingual Distributed Representations without Word Alignments. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015*, pages 748–756, Lille, France.
- Caglar Gulcehre, Sungjin Ahn, Ramesh Nallapati, Bowen Zhou, and Yoshua Bengio. 2016. Pointing the Unknown Words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 140–149, Berlin, Germany.
- Ruining He and Julian McAuley. 2016. Vbpr: Visual bayesian personalized ranking from implicit feedback. In *AAAI Conference on Artificial Intelligence*, pages 144–150. AAAI Press.
- Ann Irvine and Chris Callison-Burch. 2013. Supervised Bilingual Lexicon Induction with Multiple Monolingual Signals. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL 2013)*, pages 518–523, Atlanta, Georgia.
- Ann Irvine and Chris Callison-Burch. 2016. End-to-end statistical machine translation with zero or small parallel texts. *Natural Language Engineering*, 22(04):517–548.
- Ann Irvine and Chris Callison-Burch. 2017. A Comprehensive Analysis of Bilingual Lexicon Induction. *Computational Linguistics*, 43(2):273–310.
- Douwe Kiela, Ivan Vulić, and Stephen Clark. 2015. Visual Bilingual Lexicon Induction with Transferred ConvNet Features. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 148–158, Lisbon, Portugal.
- Alexandre Klementiev, Ann Irvine, Chris Callison-Burch, and David Yarowsky. 2012a. Toward Statistical Machine Translation without Parallel Corpora. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, pages 130–140, Avignon, France.
- Alexandre Klementiev and Dan Roth. 2006. Weakly Supervised Named Entity Transliteration and Discovery from Multilingual Comparable Corpora. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 817–824, Sydney, Australia.
- Alexandre Klementiev, Ivan Titov, and Binod Bhat-tarai. 2012b. Inducing Crosslingual Distributed Representations of Words. In *Proceedings of COLING 2012*, pages 1459–1474, Mumbai, India.
- Philipp Koehn, Franz Josef Och, and Daniel Marcu. 2003. Statistical Phrase-Based Translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, pages 48–54.
- Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer*, 42(8):30–37.
- Ang Lu, Weiran Wang, Mohit Bansal, Kevin Gimpel, and Karen Livescu. 2015. Deep Multilingual Correlation for Improved Word Embeddings. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 250–256, Denver, Colorado.

- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Bilingual Word Representations with Monolingual Quality in Mind. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 151–159, Denver, Colorado.
- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. 2013a. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013b. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems 26*, pages 3111–3119.
- Ellie Pavlick, Matt Post, Ann Irvine, Dmitry Kachaev, and Chris Callison-Burch. 2014. The language demographics of Amazon Mechanical Turk. *Transactions of the Association for Computational Linguistics*, 2:79–92.
- Reinhard Rapp. 1995. Identifying Word Translations in Non-Parallel Texts. In *Proceedings of the 33rd Annual Meeting of the Association for Computational Linguistics*, pages 320–322, Cambridge, Massachusetts.
- Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, pages 45–50, Valletta, Malta.
- Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 452–461.
- Sebastian Riedel, Limin Yao, Andrew McCallum, and Benjamin M. Marlin. 2013. Relation Extraction with Matrix Factorization and Universal Schemas. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 74–84, Atlanta, Georgia.
- Tim Rocktäschel, Sameer Singh, and Sebastian Riedel. 2015. Injecting Logical Background Knowledge into Embeddings for Relation Extraction. In *Proceedings of the 2015 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1119–1129, Denver, Colorado.
- Charles Schafer and David Yarowsky. 2002. Inducing Translation Lexicons via Diverse Similarity Measures and Bridge Languages. In *Proceedings of the 6th Conference on Natural Language Learning - Volume 20*, COLING-02, pages 1–7, Taipei, Taiwan.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Neural Machine Translation of Rare Words with Subword Units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany.
- Tianze Shi, Zhiyuan Liu, Yang Liu, and Maosong Sun. 2015. Learning Cross-lingual Word Embeddings via Matrix Co-factorization. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 567–572, Beijing, China.
- Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. 2013. Zero-Shot Learning Through Cross-Modal Transfer. In *Advances in Neural Information Processing Systems 26*, pages 935–943.
- Shyam Upadhyay, Manaal Faruqui, Chris Dyer, and Dan Roth. 2016. Cross-lingual models of word embeddings: An empirical comparison. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1661–1670, Berlin, Germany.
- Patrick Verga, David Belanger, Emma Strubell, Benjamin Roth, and Andrew McCallum. 2016. Multilingual relation extraction using compositional universal schema. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 886–896, San Diego, California.
- Ivan Vulić, Douwe Kiela, Stephen Clark, and Marie-Francine Moens. 2016. Multi-Modal Representations for Improved Bilingual Lexicon Learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 188–194, Berlin, Germany.
- Ivan Vulić and Anna Korhonen. 2016. On the Role of Seed Lexicons in Learning Bilingual Word Embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 247–257, Berlin, Germany.
- Ivan Vulić and Marie-Francine Moens. 2013. A study on bootstrapping bilingual vector spaces from non-parallel data (and nothing else). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1613–1624, Seattle, Washington.

- Ivan Vulić and Marie-Francine Moens. 2015. Bilingual Word Embeddings from Non-Parallel Document-Aligned Data Applied to Bilingual Lexicon Induction. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 719–725, Beijing, China.
- Ivan Vulic and Marie-Francine Moens. 2016. Bilingual Distributed Word Representations from Document-Aligned Comparable Data. *Journal of Artificial Intelligence Research*, 55:953–994.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *arXiv preprint arXiv:1609.08144*.
- Guangyou Zhou, Fang Liu, Yang Liu, Shizhu He, and Jun Zhao. 2013. Statistical machine translation improves question retrieval in community question answering via matrix factorization. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 852–861, Sofia, Bulgaria.
- Will Y. Zou, Richard Socher, Daniel Cer, and Christopher D. Manning. 2013. Bilingual word embeddings for phrase-based machine translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1393–1398, Seattle, Washington.