

---

# Quantized Wasserstein Procrustes Alignment of Word Embedding Spaces

Prince Osei Aboagye<sup>1\*</sup>, Yan Zheng<sup>2</sup>, Chin-Chia Michael Yeh<sup>2</sup>, Junpeng Wang<sup>2</sup>, Zhongfang Zhuang<sup>2</sup>, Huiyuan Chen<sup>2</sup>, Liang Wang<sup>2</sup>, Wei Zhang<sup>2</sup>, Jeff M. Phillips<sup>1</sup>

<sup>1</sup>University of Utah, <sup>2</sup>Visa Research

<sup>1</sup>{prince, jeffp}@cs.utah.edu

<sup>2</sup>{yazheng, miyeh, junpenwa, zzhuang, hchen, liawang, wzhan}@visa.com

---

## Abstract

Optimal Transport (OT) provides a useful geometric framework to estimate the permutation matrix under unsupervised cross-lingual word embedding (CLWE) models that pose the alignment task as a Wasserstein-Procrustes problem. However, linear programming algorithms and approximate OT solvers via Sinkhorn for computing the permutation matrix come with a significant computational burden since they scale cubically and quadratically, respectively, in the input size. This makes it slow and infeasible to compute OT distances exactly for a larger input size, resulting in a poor approximation quality of the permutation matrix and subsequently a less robust learned transfer function or mapper. This paper proposes an unsupervised projection-based CLWE model called quantized Wasserstein Procrustes (qWP). qWP relies on a quantization step of both the source and target monolingual embedding space to estimate the permutation matrix given a cheap sampling procedure. This approach substantially improves the approximation quality of empirical OT solvers given fixed computational cost. We demonstrate that qWP achieves state-of-the-art results on the Bilingual lexicon Induction (BLI) task.

## 1 Introduction

In natural language processing (NLP), the problem of aligning monolingual embedding spaces to induce a shared cross-lingual vector space has been shown not only to be useful in a variety of tasks such as bilingual lexicon induction (BLI) (Mikolov et al., 2013; Barone, 2016; Artetxe et al., 2017; Aboagye et al., 2022), machine translation (Artetxe et al., 2018b), cross-lingual information retrieval (Vulić & Moens, 2015), but it plays a crucial role in facilitating the cross-lingual transfer of language technologies from high resource languages to low resource languages.

Cross-lingual word embeddings (CLWEs) represent words from two or more languages in a shared cross-lingual vector space in which words with similar meanings obtain similar vectors regardless of their language. There has been a flurry of work dominated by the so-called *projection-based* CLWE models (Mikolov et al., 2013; Artetxe et al., 2016, 2017, 2018a; Smith et al., 2017; Ruder et al., 2019), which aim to improve CLWE model performance significantly. *Projection-based* CLWE models learn a transfer function or mapper between two independently trained monolingual word vector spaces with limited or no cross-lingual supervision.

Famous among *projection-based* CLWE models are the unsupervised *projection-based* CLWE models (Artetxe et al., 2017; Lample et al., 2018; Alvarez-Melis & Jaakkola, 2018;

---

\*work done while interning at Visa Research

Grave et al., 2019): they eliminate the initial seed bilingual lexicon and rely on the topological similarities between monolingual spaces, known as the isometry assumption, to extract seed bilingual lexicons. This makes them attractive since they require no cross-lingual supervision. One of the ways of framing unsupervised CLWE models is to pose the alignment task as a Wasserstein-Procrustes problem aiming to jointly estimate a permutation matrix and an orthogonal matrix (Grave et al., 2019; Ramírez et al., 2020). Most existing unsupervised CLWE models that solve the Wasserstein-Procrustes problem resort to Optimal Transport (OT) based methods to estimate the permutation matrix.

Optimal Transport (OT) (Monge, 1781; Kantorovich, 1942) provides a natural geometric and probabilistic toolbox to compare probability distributions or measures. OT is concerned about determining an optimal transport plan for moving probability mass between two probability distributions with the cheapest cost. In theory, optimal transport is beautiful and well defined and has been well studied under continuous distribution. However, in practice or specifically in machine learning, we only have access to samples given an underlying distribution, so we turn to observe discrete distributions. This resonates with how empirical OT solvers have been built; they accept samples as inputs from input probability distributions or measures.

When the discrete distributions are composed of a large number of point cloud in higher dimensions, it becomes slow, impractical, and infeasible to compute OT distances exactly given the empirical OT solvers. A common scalable approach adopted by Grave et al. (2019) in their stochastic optimization framework to approximate the exact OT distance in order to extract the permutation matrix was to randomly draw  $k$  monolingual embeddings from the source and target spaces, respectively. However, this approximation approach poses two main challenges:

**1) Sampling Efficiency** Does the OT distance computed between the  $k$  sampled embeddings provide a useful or quality OT distance approximation of the true underlying distributions of the source and target spaces? Theoretical bounds and results have shown that the quality of this approximation has a convergence rate of  $k^{-\frac{1}{d}}$  to the true OT distance, where  $d$  is the ambient dimension (Dudley, 1969; Weed & Bach, 2019). Therefore, an effective approximation of the true OT distance requires large  $k$  samples since we are constrained by the curse of dimensionality from the power  $-\frac{1}{d}$ . Thus, we need more samples to approximate the true OT distance in higher dimensions.

**2) Computational Efficiency** Empirical OT solvers such as linear programming algorithms (Burkard et al., 2012) and approximate solvers via Sinkhorn (Cuturi, 2013) for computing the permutation matrix have a computational cost of  $\mathcal{O}(k^3 \log k)$  and  $\mathcal{O}(k^2 \epsilon^{-2})$ , respectively, in the input size,  $k$ , and regularization term  $\epsilon$  defined later in Equation 7. It becomes slow and infeasible in higher dimensions to compute OT distances exactly for a larger input size. We are therefore restricted by the maximum  $k$  samples to draw for an effective approximation of the true OT distance. The constraint here is not the availability of data but computational cost.

Given these two challenges, Beugnot et al. (2021) proposed two efficient OT estimators. The empirical OT solvers remain the same, either the linear programming solver or the entropic-regularized OT via Sinkhorn. However, instead of drawing only  $k$  samples as input to the OT solver, they rely on a cheap quantization step like  $k$ -means ++ (Arthur & Vassilvitskii, 2007) that is consistent with the computational complexity of the OT solver. Since sampling is cheap, they draw more than  $k$  samples and then use  $k$ -means++ to quantize the oversampled points from the source and target spaces, respectively, by partitioning them into  $k$  clusters and then select the  $k$  weighted anchor points as input to the OT solver. This quantization step improves the approximation quality to the true OT distance. Aside from the theoretical guarantees of the benefits of this quantization step, they showed that the new variant of the unregularized OT

estimator yield an improvement in the convergence rate by  $k^{-2\alpha}$  in the best case or  $k^{-\alpha}$  in the worst case, which is on par with the computational complexity existing empirical OT estimators, where  $\alpha = \frac{1}{d}$ .

Inspired by the work of [Beugnot et al. \(2021\)](#), our paper proposes a new unsupervised CLWE model called quantized Wasserstein Procrustes (qWP). We follow the stochastic algorithm framework by [Grave et al. \(2019\)](#) and the refinement procedure from [Lample et al. \(2018\)](#).

**Our contribution.** This work proposes a new unsupervised CLWE model: **quantized Wasserstein Procrustes (qWP)** that relies on a quantization step of the source and target distributions to estimate the alignment and linear transformation jointly. Firstly, we use the stochastic optimization framework in [Grave et al. \(2019\)](#). However, instead of randomly drawing  $k$  samples at each iteration, we use a quantization step to preprocess the source and target distributions to find the optimal  $k$  point compression or summary needed to estimate the permutation matrix. It leads to a much-refined sample as opposed to a random sampling of the  $k$  points. This approach substantially improves the approximation quality of the true OT distance and bias of empirical OT solvers given fixed computational cost ([Beugnot et al., 2021](#)). The main idea behind qWP is to oversample the  $k$  samples and then reduce them to  $k$ -weighted samples through quantization such as  $k$ -means++. After this, a linear program solver or regularized Sinkhorn algorithm can be used on the resulting quantized distribution. The translation pairs obtained from the permutation matrix are then used to learn the linear transformation. Finally, we use the refinement approach from [Lample et al. \(2018\)](#) to improve the orthogonal mapping. We demonstrate that qWP achieves state-of-the-art results on the BLI task.

## 2 Related Work

At the heart of Cross-lingual NLP are CLWE models. It has quickly evolved into a large subarea with a wide variety of approaches and perspectives, so we provide context by overviewing this work first.

Projection-based CLWE models can be categorized into ([Ruder et al., 2019](#)): **1**) fully supervised projection-based CLWE models, **2**) weakly supervised projection-based CLWE models, and **3**) fully unsupervised projection-based CLWE models. The main idea governing all CLWE models is to independently train monolingual embeddings on large monolingual corpora in different languages or use pre-trained monolingual embeddings and then learn a transfer function to map them into a shared cross-lingual word vector space.

The first fully supervised projection-based CLWE model to learn a shared cross-lingual word vector space from monolingually-trained word embedding was proposed by [Mikolov et al. \(2013\)](#). They learned a linear transform from the source embedding space to the target language by minimizing the sum of squared Euclidean distance between the translation pairs of a seed dictionary based on the assumption that two embedding spaces exhibit similar geometric structures (i.e., approximately isomorphic). Their model requires word-level supervision from several thousand seed translation dictionaries (Dict). Subsequent works by [Xing et al. \(2015\)](#); [Artetxe et al. \(2016\)](#); [Smith et al. \(2017\)](#) argued and proved that the quality of the learned CLWEs could be improved by modifying the objective function in [Mikolov et al. \(2013\)](#).

A more recent line of research has shown that the shared cross-lingual word vector space can be induced with weaker supervision from a small initial seed dictionary ([Vulic & Korhonen, 2016](#); [Glavaš et al., 2019](#); [Vulić et al., 2019](#)). Weakly supervised projection-based CLWE models start with a small initial seed dictionary; however, the initial seed dictionary is iteratively expanded through a self-learning procedure. For example, Bootstrap Procrustes (PROC-B) ([Glavaš et al., 2019](#)) is semi-supervised in that it starts with a small pairwise correspondence (of 500-1000 words), aligns those to infer a larger correspondence, and repeats applying Procrustes alignment. The quest to eliminate cross-lingual supervision has led to the development of fully unsupervised

projection-based CLWE models.

Fully unsupervised projection-based CLWE models use the topological similarities between monolingual embedding spaces to induce the shared cross-lingual vector space (Lample et al., 2018; Artetxe et al., 2018a; Mohiuddin & Joty, 2019). The translation dictionaries are produced from scratch based on monolingual data only.

### 3 Background

In this section, we describe the mathematical formulation of supervised *projection-based* CLWE models and unsupervised *projection-based* CLWE models. We also defined what the 2-Wasserstein distance is and looked in detail at how the Wasserstein-Procrustes problem under the unsupervised CLWE model is solved in practice.

We define two monolingual embedding spaces as  $X, Y \in \mathbb{R}^{n \times d}$ , where  $n$  is the number of words, and  $d$  is the dimension of the monolingual word embeddings.

**Supervised Projection-Based CLWE Models** require word-level supervision from seed translation dictionaries such that word  $x_i$  in  $X$  is the translation of word  $y_i$  in  $Y$ . The linear transformation,  $W^*$ , from the source monolingual embedding space to the target monolingual embedding space is learned by solving the least square problem (Mikolov et al., 2013):

$$W^* = \arg \min_{W \in \mathbb{R}^{d \times d}} \|XW - Y\|_F^2 \quad (1)$$

Xing et al. (2015), modified the objective function in Eq. (1) to improve the quality of the learned CLWEs by unit length normalizing the word embeddings and imposing an orthogonality constraint on the linear transformation ( $W$ ) during training:

$$W^* = \arg \min_{W \in \mathcal{O}_d} \|XW - Y\|_F^2, \quad (2)$$

where  $\mathcal{O}_d$  is the set of orthogonal matrices. The orthogonality constraint preserves the original monolingual embedding space’s similarities and geometric structure. These assumptions and constraints imposed on the linear transform make the problem of learning a transfer function an orthogonal **Procrustes** problem (Eq. 2), which has a closed-form solution:  $W^* = UV^\top$ , where  $U\Sigma V^\top$  is the singular value decomposition of  $X^\top Y$  (Schönemann, 1966).

**2-Wasserstein distance** is a distance function used to compute the OT-distance given two set of points  $X$  and  $Y$ :

$$W_2^2(X, Y) = \min_{P \in \mathcal{P}_n} \sum_{i,j=1}^n \|x_i - y_j\|_2^2 P_{ij} \quad (3)$$

where  $\mathcal{P}_n$  is the set of permutation matrices,  $\mathcal{P}_n = \{P \in \{0, 1\}^{n \times n}, P1_n = 1_n, P^\top 1_n = 1_n\}$ .

**Unsupervised Projection-Based CLWE Models** Without any initial seed bilingual lexicon some unsupervised CLWE models solves the **Wasserstein-Procrustes** problem (Eq. 4) to jointly estimate the permutation matrix or alignment ( $P$ ) and linear transformation ( $W$ ) (Grave et al., 2019; Ramírez et al., 2020):

$$W^*, P^* = \arg \min_{W \in \mathcal{O}_d, P \in \mathcal{P}_n} \|XW - PY\|_F^2 \quad (4)$$

The permutation matrix  $P^*$  provides a one-to-one mapping or correspondence between the source and target samples.

Under unsupervised CLWE models that solve the Wasserstein-Procrustes problem, we aim to estimate the two unknown variables  $W$  and  $P$ . One way to solve Eq. (4) is by alternating the minimization of  $W$  and  $P$ . Given  $P$ , we use the translation pairs obtained between the source and target spaces to learn the linear transformation,  $W^*$  from Eq. (2). Similarly, given the linear transformation  $W^*$ , Eq. (4) is equivalent to minimizing the 2-Wasserstein distance between  $XW$  and  $Y$  to solve for the permutation matrix,  $P$ :

$$W_2^2(XW, Y) = \min_{P \in \mathcal{P}_n} \sum_{i,j=1}^n \|x_i W - y_j\|_2^2 P_{ij} \quad (5)$$

Equation (5) is the standard OT problem, and it can be solved using a linear programming solver, which has a computational cost of  $\mathcal{O}(n^3 \log n)$ . For a large  $n$ , a linear programming solver is impractical. Another variant and approximation of the optimal transport problem were proposed by (Cuturi, 2013). This variant adds an entropic regularization term leading to the Sinkhorn algorithm with a computational cost of  $\mathcal{O}(n^2 \epsilon^{-2})$ :

$$W_2^2(XW, Y) = \min_{P \in \mathcal{P}_n} \sum_{i,j=1}^n \|x_i W - y_j\|_2^2 P_{ij} + \epsilon \sum_{i,j=1}^n \log P_{ij} \quad (6)$$

Grave et al. (2019) proposed a stochastic optimization scheme to jointly estimate  $W$  and  $P$  by randomly sampling  $\hat{X}, \hat{Y} \in \mathbb{R}^{k \times d}$  from  $X$  and  $Y$ , where  $k < n$ . Due to how slow and infeasible a linear programming solver for a larger input size can be, Grave et al. (2019) used the Sinkhorn algorithm to compute the permutation matrix,  $P$  by minimizing:

$$W_2^2(\hat{X}W, \hat{Y}) = \min_{P \in \mathcal{P}_k} \sum_{i,j=1}^k \|x_i W - y_j\|_2^2 P_{ij} + \epsilon \sum_{i,j=1}^k \log P_{ij} \quad (7)$$

## 4 Proposed Method

This section introduces our new unsupervised CLWE model: quantized Wasserstein Procrustes (qWP). We use the previous stochastic algorithm framework and refinement procedure from Grave et al. (2019) and Lample et al. (2018) respectively in our model, but we rely on a quantization step to estimate the permutation matrix.

### 4.1 quantized Wasserstein Procrustes (qWP)

We consider two languages with vocabularies  $V_x$  and  $V_y$ , represented by word embeddings  $X = \{x_i\}_{i=1}^n, Y = \{y_i\}_{i=1}^n$ , respectively. We assume two empirical distributions over the embedding spaces,  $X$  and  $Y$ :  $\mu = \sum_{i=1}^n p_i \delta_{x^{(i)}}$  and  $\nu = \sum_{j=1}^n q_j \delta_{y^{(j)}}$ , where  $p_i$  and  $q_i$  are the probability weights associated with each word vector,  $\delta_x$  and  $\delta_y$  is the Dirac function supported on point  $x$  and  $y$  respectively.

The main crux of our proposed unsupervised CLWE model: quantized Wasserstein Procrustes (qWP) is that we rely on a quantization step like  $k$ -means++ (Arthur & Vassilvitskii, 2007) instead of random sampling to estimate the permutation matrix and then use gradient descent and Procrustes to extract the orthogonal matrix. We take Eq. (4) as our loss function. However, Eq. (4) is not jointly convex in  $W$  and  $P$ , but as we saw in Section 3 we can fix one variable and then solve for the other variable. Alternating the minimization in each variable  $W$  and  $P$  is therefore employed to find a solution (Alaux et al., 2018; Grave et al., 2019).

First, we have to induce the translation dictionary by solving for the permutation matrix,  $P^*$  in Eq. (5) and then find the orthogonal projection matrix from Eq. (2). Naively doing an

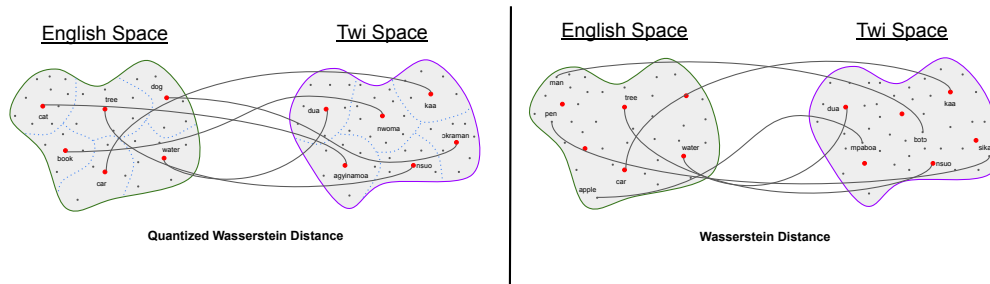


Figure 1: Illustration on toy  $2d$  data showing the potential advantage of Quantized Wasserstein Distance (qWD) over Wasserstein Distance (WD). We want to align or translate words in the English Space to words in the Twi Space without knowing beforehand the translation pairs or the linear transformation. Twi is a language spoken in Ghana, West Africa. First, we must induce the translation pairs by estimating the permutation matrix,  $P$ , either through qWD or WD. Each dot represents a word in that space; specifically, the red points are the  $k$  centers from  $k$ -means++. The edge connecting two red points means the two words are accurate translation pairs, whereas the edge between two black points is the wrong translation pair. Here we want to induce six translation pairs through  $P$ .

alternating full minimization in each variable  $W$  and  $P$  of Eq. (4) does not scale, and even on smaller problems, empirical results show that it quickly converges to a bad local minima (Zhang et al., 2017). A scalable stochastic approach adopted by Grave et al. (2019) was to instead, at each iteration,  $t$ , randomly sample a minibatch  $X_k = \{x_i\}_{i=1}^k$ , and  $Y_k = \{y_i\}_{i=1}^k$  of size  $k$  from  $X$  and  $Y$ . The optimal coupling or permutation matrix,  $P^*$ , was then computed from Eq. (7) using the Sinkhorn algorithm. The translation pairs obtained from  $P^*$  between the source and target spaces are then used to learn the orthogonal matrix,  $W^*$ , that maps the source to the target spaces from Eq. (2) by using Procrustes and gradient descent to update  $W$ . The procedure for updating  $W$  is detailed in Grave et al. (2019).

The stochastic optimization scheme adopted by Grave et al. (2019) to make the alternating minimization process scale and achieve a better convergence to a good local minimum when computing the permutation matrix suffers from the sampling efficiency and computational efficiency challenges discussed in Section 1.

To address these two challenges following Beugnot et al. (2021), we will quantize the source and target word embedding space by finding the optimal  $k$  point compression or summary as input to the 2-Wasserstein distance (Pollard, 1982; Canas & Rosasco, 2012) through the use of  $k$ -means++. The resulting convergence rate of  $k^{-2\alpha}$  in the best case or  $k^{-\alpha}$  in the worst case from using this quantization step makes the OT solver yields a better approximation quality of the permutation matrix and subsequently a more robust learned transfer function, where  $\alpha = \frac{1}{d}$ .

#### 4.1.1 New Alignment Algorithm

The goal of our proposed new CLWE algorithm is to quantize the source and target embedding spaces  $X$  and  $Y$  to be aligned to obtain a much-refined coreset<sup>1</sup> that is less noisy compared to just randomly sampling from  $X$  and  $Y$ . Our proposed new method is summarized in Algorithms 1 and 2. For each iteration  $t$  (Algorithm 1), we compute the permutation matrix  $P^*$  from Algorithm 2. The main idea of Algorithm 2 is to draw more than  $k$  samples using the coreset size  $m > k$  and then reduce them to  $k$ -weighted samples through quantization such as  $k$ -means++. Here the computational cost of  $k$ -means++ is  $\mathcal{O}(mk)$ . To satisfy the computational complexity of the OT

<sup>1</sup>A coreset is a summary or an approximation of the shape of a larger point cloud with a smaller point cloud.



---

**Algorithm 1** Quantized Wasserstein Procrustes

---

**Require:**

Word embedding matrix,  $X, Y \in \mathbb{R}^{n \times d}$  of the source language and target language respectively, entropy regularization coefficient  $\epsilon$ , number of anchor point  $k$

**Ensure:**

Orthogonal matrix,  $W$

- 1: **for**  $e = 1, \dots, E$  **do**
  - 2:   **for**  $t = 1, \dots, T$  **do**
  - 3:      $P \leftarrow qW(X, Y, \epsilon, k)$
  - 4:      $W \leftarrow$  Update  $W$  by gradient descent and Procrustes
  - 5:   **end for**
  - 6: **end for**
  - 7: **return**  $W$
- 

solver, we must ensure that the quantization step used to preprocess the source and target space takes  $\mathcal{O}(k^3 \log k)$  time. In view of this, we set  $m = k^2 \log k$  so that we are consistent with the computational complexity  $\mathcal{O}(k^3 \log k)$  of the OT solver. We then sample  $X_m = (x_1, \dots, x_m)$  i.i.d from  $X$  and  $Y_m = (y_1, \dots, y_m)$  i.i.d from  $Y$ . Using  $k$ -means++ we find the  $k$  weighted centers. Following each Voronoi cell, we weight each center proportionally to the number of samples to obtain the weights  $a$  and  $b$ . We then can use either the linear program solver or the regularized Sinkhorn algorithm (Cuturi, 2013) to estimate the permutation matrix,  $P$ , between the two quantized point clouds. In our case, we used the entropic-regularized OT solver via Sinkhorn, which we call APPROXOT( $C, a, b, \epsilon$ ).

---

**Algorithm 2** Quantized 2-Wasserstein Distance ( $qW(X, Y, \epsilon, k)$ )

---

**Require:**

$X = \{x_i\}_{i=1}^n, Y = \{y_i\}_{i=1}^n$ , entropy regularization coefficient  $\epsilon$ , number of anchor points  $k$

**Ensure:**

Permutation Matrix,  $P$

- 1: Sample  $m$  points:
  - 2:   Set  $m = k^2 \log k$
  - 3:   Sample  $X_m = (x_1, \dots, x_m)$  i.i.d from  $X$  and  $Y_m = (y_1, \dots, y_m)$  i.i.d from  $Y$
  - 4: Subsample  $k$  anchor points:
  - 5:   Compute  $(c_1, \dots, c_k)$  with  $k$ -means++
  - 6:   Compute  $(d_1, \dots, d_k)$  with  $k$ -means++
  - 7: Compute weights:
  - 8:   Set  $a_i = \sum_{j=1}^n \mathbf{1}_{i=\arg \min_t \|x_j - c_t\|_2} \forall i \in \{1, \dots, k\}$
  - 9:   Set  $b_i = \sum_{j=1}^n \mathbf{1}_{i=\arg \min_t \|x_j - d_t\|_2} \forall i \in \{1, \dots, k\}$
  - 10: Cost matrix:
  - 11:   Set  $C_{ij} = \|c_i - d_j\|_2^2 \forall i, j \in \{1, \dots, n\}$
  - 12: Regularized transport solver:
  - 13: **return**  $P \leftarrow$  APPROXOT( $C, a, b, \epsilon$ )
- 

See the example in Figure 1 where the translation pairs obtained under qWD yield perfect matches compared to WD, which gave some wrong translation pairs. Under qWD we use  $k$ -means++ to quantize the English and Twi Space to select the  $k$  weighted centers as input to

the OT solver instead of randomly drawing  $k$  points under WD, which could be noisy.

As a quick review of  $k$ -means++ (Arthur & Vassilvitskii, 2007), it initializes a set of cluster centers for the  $k$ -means objective. Each step iteratively increases the set of cluster centers by choosing a new center from the dataset proportional to the squared distance to the closest already chosen center. In one variant we explore, we run one step of the standard Lloyd’s algorithm after initializing, moving each center found to the average of data points closest to it.

## 5 Experimental Analysis

We provide an evaluation of our proposed methods using English (EN) and five languages embeddings pre-trained on Wikipedia (Bojanowski et al., 2017): Spanish (ES), French (FR), German (DE), Russian (RU), and Italian (IT). We use the 300-dimensional fastText (Bojanowski et al., 2017) embeddings, and all vocabularies are trimmed to the 200K most frequent words.

**Alignment evaluation tasks: BLI** We evaluate and compare our proposed CLWE method mainly on the Bilingual Lexicon Induction (BLI) task, a word translation task. BLI is more direct and has become the de facto evaluation task for CLWE models. For words in the source language, this task retrieves the nearest neighbors in the target language after alignment to check if it contains the translation. We report two different translation accuracies: precision at 1 (P@1) and mean average precision (MAP) (Glavaš et al., 2019) translation accuracy, which is equivalent to the mean reciprocal rank (MRR) of the translation.

**Implementation Details** The monolingual word embeddings are unit length normalized and centered before entering the model. The first 2.5k words are used to determine  $Q_0$  given  $P^*$  obtained from the Frank-Wolfe algorithm (Frank & Wolfe, 1956). We trained qWP on the first 20k most frequent words and evaluated them on separate 1.5k source test queries. We used the MUSE publicly available translation dictionary (Lample et al., 2018). We used the regularized Sinkhorn algorithm (Cuturi, 2013) and always set the entropy regularization term ( $\epsilon$ ) to  $\epsilon = 0.05$ .

We use the *Refinement* approach from (Lample et al., 2018) and run it for five epochs. This approach iteratively improves the orthogonal mapping  $Q$ . After learning  $Q^*$  from Eq. (4), we build another (slightly larger) dictionary of translation pairs by translating each word to its nearest neighbor under the transformation  $Q$ . The newly learned dictionary of translation pairs is then used to learn a new mapping  $Q$  from Eq. (2), and then we repeat the process, each time building an incrementally larger dictionary.

We consider both balanced and unbalanced OT. The unbalanced OT does not require strict mass preservation (Chizat et al., 2018), contrary to the standard or balanced OT problem, Eq. (5). Under the unbalanced OT, Eq. (5) is relaxed by adding two KL-divergence terms to ensure a more relaxed mass preservation. This helps to solve the polysemy problem.

**Baselines: BLI** We evaluated and compared the published result of qWP to several supervised and unsupervised CLWE models on the BLI task. The baselines include Procrustes (PROC) (Artetxe et al., 2016), Ranking-Based Optimization (RCSLS) (Joulin et al., 2018), Gromov Wasserstein (GW) (Alvarez-Melis & Jaakkola, 2018), Adversarial Training (Adv + Refine) (Lample et al., 2018) and the density matching method (Dema + Refine) (Wang et al., 2019). We used the baseline results

**Main Results** Tables 1, 2 and 3 summarize the effect of the coreset size within the qWP algorithm. We proceed with four experiments. In tables 1 and 3 we report the mean average precision (MAP) (Glavaš et al., 2019) translation accuracy, which is equivalent to the mean reciprocal rank (MRR) of the translation, whereas, Tables 2 and 4, the translation accuracy reported is the precision at 1 (P@1).



Table 1: Bilingual lexicon Induction (BLI) task, (MAP) - Without Refinement

Trans. Pairs	Sampling	Coreset Size				
		200	500	1000	2000	3000
EN-ES	Random	36.40	47.22	48.90	49.66	50.08
	KMeans ++	45.21	48.69	49.64	49.71	50.09
ES-EN	Random	43.74	50.36	52.04	53.84	54.67
	KMeans ++	47.24	52.21	52.55	54.10	54.90
EN-FR	Random	37.22	47.94	49.31	50.59	50.88
	KMeans ++	46.54	49.45	50.12	50.54	51.07
FR-EN	Random	38.87	53.36	54.91	55.47	55.85
	KMeans ++	52.67	54.43	55.54	56.11	56.70
EN-DE	Random	27.18	36.10	38.00	38.49	39.29
	KMeans++	32.80	37.44	38.58	38.77	39.72
DE-EN	Random	30.97	41.26	40.44	41.87	41.78
	KMeans ++	39.03	41.90	40.21	43.93	42.42
EN-RU	Random	18.68	27.91	30.91	31.75	32.43
	KMeans ++	26.97	27.12	29.90	32.25	31.41
RU-EN	Random	27.26	39.93	41.50	43.56	43.81
	KMeans ++	16.13	37.07	42.69	42.82	44.54
EN-IT	Random	34.04	46.10	47.99	49.28	50.79
	KMeans ++	44.83	47.31	49.00	50.29	51.12
IT-EN	Random	38.50	52.44	52.92	54.70	57.04
	KMeans ++	47.80	51.77	54.60	57.03	57.49
Avg	Random	33.28	44.26	45.69	46.92	47.66
	KMeans ++	<b>39.92</b>	<b>44.74</b>	<b>46.28</b>	<b>47.55</b>	<b>47.94</b>

Table 2: Bilingual lexicon Induction (BLI) task, (P@1) Without Refinement

Translation Pairs	Sampling	Coreset Size			
		500	1000	2000	3000
EN-ES	Random	73.53	75.20	76.73	80.40
	KMeans ++	77.80	79.47	78.20	81.53
EN-FR	Random	77.07	79.40	80.00	81.00
	KMeans ++	78.27	79.60	80.20	81.13
EN-DE	Random	62.73	67.60	70.40	70.60
	KMeans++	65.60	68.87	71.40	71.40
EN-RU	Random	33.13	35.53	35.47	36.87
	KMeans ++	34.53	36.07	36.53	36.60
EN-IT	Random	70.67	72.47	75.13	75.73
	KMeans ++	73.20	74.87	76.73	76.93
Avg	Random	63.43	66.04	67.55	68.92
	KMeans ++	<b>65.88</b>	<b>67.78</b>	<b>68.61</b>	<b>69.52</b>

The first experiments in Table 1 show the MRR scores without refinement, and the following Table 3 shows the same MRR scores with refinement. In each table, we increase the coreset size from 200 to 3000, and this is either chosen as in prior work as a random sample or in our proposed approach via  $k$ -means++. As expected, on all language pairs, the performance increases as the coreset size increases. Also, notice that the improvement by increasing the coreset size plateaus and is not as significant from 2000 to 3000, indicating that probably 2000 coreset points are usually sufficient.

We also observe that in almost all cases, the performance is improved when using the  $k$ -means++ coreset instead of the random sample coreset. The few exceptions are mostly in the comparison with Russian (RU) with refinement, but this gap narrows as the coreset size increases. Notably, by coreset size of 2000, the  $k$ -means++ coresets have a clear advantage with an average improvement of from 46.92 to 47.55 without refinement and from 53.05 to 53.76 with refinement. This follows the general trend of better scores when the refinement phase is used.

Table 2 shows a similar experiment on the BLI tasks but reports the precision at 1 (P@1) score. The results show a strong average improvement while using  $k$ -means++, with the exception being EN-RU with a small advantage of random sampling at 3000 coreset size; however, with MAP, the results for  $k$ -means++ are already basically as good with 2000 points.

Table 3: Bilingual lexicon Induction (BLI) task, (MAP) With Refinement

Trans. Pairs	Sampling	Coreset Size				
		200	500	1000	2000	3000
EN-ES	Random	54.45	54.35	54.54	54.56	54.61
	KMeans ++	54.41	54.48	54.55	54.67	54.72
ES-EN	Random	60.96	58.24	58.56	58.88	59.69
	KMeans ++	58.01	58.26	59.22	59.11	59.55
EN-FR	Random	54.93	55.26	55.31	55.31	55.24
	KMeans ++	55.05	55.41	55.44	55.38	55.30
FR-EN	Random	56.00	61.36	61.44	61.46	61.51
	KMeans ++	61.81	61.68	61.54	61.60	61.64
EN-DE	Random	43.42	43.28	43.42	43.46	43.37
	KMeans++	43.12	43.32	43.56	43.59	43.52
DE-EN	Random	48.45	48.70	45.74	46.03	46.72
	KMeans ++	45.91	49.05	46.69	48.78	48.54
EN-RU	Random	40.34	41.56	42.92	42.50	42.76
	KMeans ++	41.57	40.08	41.41	43.07	41.39
RU-EN	Random	48.01	49.28	48.64	50.09	50.48
	KMeans ++	38.69	46.24	50.16	49.05	50.43
EN-IT	Random	55.93	56.82	57.36	57.48	57.54
	KMeans ++	56.23	56.55	57.32	57.75	57.41
IT-EN	Random	59.71	61.44	60.22	60.70	65.10
	KMeans ++	60.13	59.55	60.61	64.62	64.71
Avg	Random	<b>52.22</b>	<b>53.02</b>	52.82	53.05	53.70
	KMeans ++	51.49	52.46	<b>53.05</b>	<b>53.76</b>	<b>53.72</b>

The final experiment in Table 4 shows the results of our proposed methods against state-of-

the-art techniques. We used a fixed coreset size of 2000. Each entry shows the P@1 scores on the BLI task. The first two lines show PROC and RCSLS, which are supervised methods, so they know the alignment between 5000 pairs of works across embeddings and use this knowledge to determine the alignment. Notice our techniques (which are unsupervised) improve upon the standard Procrustes alignment (PROC) and are almost competitive with the RCSLS method, which optimizes for the BLI task specifically.

Our method also outperforms Gromov-Wasserstein (GW) alignment, as well as Adv + Refine, Dema + Refine, and a random sample coreset when using refinement.

In this table, we also show experiments with two other enhancements. The first is to improve the cluster centers and the quantization found with  $k$ -means++ with a run of Lloyd’s algorithm (the standard  $k$ -means optimization procedure) for 1 step. This moves the quantization point to the center of the points it represents, making it more representative on average. This provides a small improvement. The second extension is to use unbalanced optimal transport instead of balanced OT. Surprisingly, this offers no advantage on average.

Table 4: Bilingual lexicon Induction (BLI) task, Comparison with other Methods

Method	Dict	EN-ES		EN-FR		EN-DE		EN-RU		EN-IT		Avg
		→	←	→	←	→	←	→	←	→	←	
PROC	5K	81.9	83.4	82.1	82.4	74.2	72.7	51.7	63.7	77.4	77.9	74.7
RCSLS	5K	84.1	86.3	83.3	84.1	79.1	76.3	57.9	67.2			77.3
GW	None	81.7	80.4	81.3	78.9	71.9	78.2	45.1	43.7	78.9	75.2	71.5
Adv + Refine	None	81.7	83.3	82.3	82.1	74.0	72.2	44.0	59.1	77.9	77.5	73.4
Dema + Refine	None	82.8	84.9	82.6	82.4	75.3	74.9	46.9	62.4			74.0
Random												
WP + Refine	None	82.8	84.1	82.6	82.9	75.4	73.3	43.7	59.1			73.0
<b>Unbalanced OT</b>												
(Ours) KMeans++												
qWP + Refine	None	83.9	84.5	83.6	83.1	77.0	74.9	48.0	60.1	80.5	80.7	<b>75.6</b>
(Ours) LloydRefine												
qWP + Refine	None	83.8	84.9	84.3	83.4	77.0	75.2	48.2	61.3	80.5	80.9	<b>75.9</b>
<b>Balanced OT</b>												
(Ours) KMeans++												
qWP + Refine	None	83.5	84.3	84.0	83.1	76.9	74.9	46.6	59.8	80.6	80.3	<b>75.4</b>
(Ours) LloydRefine												
qWP + Refine	None	83.6	84.4	84.0	83.1	77.1	74.8	47.3	60.4	80.1	80.4	<b>75.5</b>

## 6 Conclusion

This paper presents an approach to aligning embeddings in high-dimensional space. While the overall problem is non-convex and computationally expensive, we present an efficient stochastic algorithm to solve the problem based on a refined sample set. This paper focuses on the matching procedure of the BLI task. Our key insight is that our quantization algorithm can outperform the current state-of-art unsupervised algorithm on both balanced and unbalanced settings of the loss function.

## References

- Prince Osei Aboagye, Jeff Phillips, Yan Zheng, Junpeng Wang, Chin-Chia Michael Yeh, Wei Zhang, Liang Wang, and Hao Yang. Normalization of language embeddings for cross-lingual alignment. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=Nh7CtbyoqV5>.
- Jean Alaux, Edouard Grave, Marco Cuturi, and Armand Joulin. Unsupervised hyper-alignment for multilingual word embeddings. In *International Conference on Learning Representations*, 2018.
- David Alvarez-Melis and Tommi S. Jaakkola. Gromov-wasserstein alignment of word embedding spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 1881–1890, 2018.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. Learning principled bilingual mappings of word embeddings while preserving monolingual invariance. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2289–2294, Austin, Texas, November 2016. Association for Computational Linguistics. doi: 10.18653/v1/D16-1250. URL <https://aclanthology.org/D16-1250>.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 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)*, pp. 451–462, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1042. URL <https://aclanthology.org/P17-1042>.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 789–798, Melbourne, Australia, July 2018a. Association for Computational Linguistics. doi: 10.18653/v1/P18-1073. <https://github.com/artetxem/vecmap>.
- Mikel Artetxe, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. Unsupervised neural machine translation. In *Proceedings of the Sixth International Conference on Learning Representations*, April 2018b.
- David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *ACM Symposium on Discrete Algorithms*, 2007.
- Antonio Valerio Miceli Barone. Towards cross-lingual distributed representations without parallel text trained with adversarial autoencoders. In *Proceedings of the 1st Workshop on Representation Learning for NLP*, pp. 121–126, 2016.
- Gaspard Beugnot, Aude Genevay, Justin M Solomon, and Kristjan Greenewald. Improving approximate optimal transport distances using quantization. In *UAI 2021: Uncertainty in Artificial Intelligence*, 2021.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017. doi: 10.1162/tacl\_a\_00051. URL <https://www.aclweb.org/anthology/Q17-1010>.

- Rainer Burkard, Mauro Dell’Amico, and Silvano Martello. *Assignment Problems. Revised reprint*. SIAM - Society of Industrial and Applied Mathematics, 2012. ISBN 978-1-611972-22-1. 393 Seiten.
- Guillermo Canas and Lorenzo Rosasco. Learning probability measures with respect to optimal transport metrics. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL <https://proceedings.neurips.cc/paper/2012/file/c54e7837e0cd0ced286cb5995327d1ab-Paper.pdf>.
- Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. An interpolating distance between optimal transport and fisher-rao metrics. *Foundations of Computational Mathematics*, 18(1):1–44, 2018.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013. URL <https://proceedings.neurips.cc/paper/2013/file/af21d0c97db2e27e13572cbf59eb343d-Paper.pdf>.
- R. M. Dudley. The speed of mean glivenko-cantelli convergence. *Annals of Mathematical Statistics*, 40(1):40–50, 1969.
- Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956. doi: <https://doi.org/10.1002/nav.3800030109>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/nav.3800030109>.
- Goran Glavaš, Robert Litschko, Sebastian Ruder, and Ivan Vulić. How to (properly) evaluate cross-lingual word embeddings: On strong baselines, comparative analyses, and some misconceptions. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 710–721, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1070. URL <https://aclanthology.org/P19-1070>.
- Edouard Grave, Armand Joulin, and Quentin Berthet. Unsupervised alignment of embeddings with wasserstein procrustes. In Kamalika Chaudhuri and Masashi Sugiyama (eds.), *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1880–1890. PMLR, 16–18 Apr 2019. URL <https://proceedings.mlr.press/v89/grave19a.html>.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. Loss in translation: Learning bilingual word mapping with a retrieval criterion. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2979–2984, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1330. <https://github.com/facebookresearch/fastText/tree/master/alignment>.
- Leonid V Kantorovich. On the translocation of masses. In *Dokl. Akad. Nauk. USSR (NS)*, volume 37, pp. 199–201, 1942.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. In *International Conference on Learning Representations*, 2018.

- Tomas Mikolov, Quoc V Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*, 2013.
- Tasnim Mohiuddin and Shafiq Joty. Revisiting adversarial autoencoder for unsupervised word translation with cycle consistency and improved training. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 3857–3867, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1386. URL <https://aclanthology.org/N19-1386>.
- Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale Sci.*, pp. 666–704, 1781.
- D. Pollard. Quantization and the method of k-means. *IEEE Transactions on Information Theory*, 28(2):199–205, 1982. doi: 10.1109/TIT.1982.1056481.
- Guillem Ramírez, Rumén Dangovski, Preslav Nakov, and Marin Soljačić. On a novel application of wasserstein-procrustes for unsupervised cross-lingual learning. *arXiv preprint arXiv:2007.09456*, 2020.
- Sebastian Ruder, Ivan Vulić, and Anders Søgaard. A survey of cross-lingual word embedding models. *J. Artif. Int. Res.*, 65(1):569–630, may 2019. ISSN 1076-9757. doi: 10.1613/jair.1.11640. URL <https://doi.org/10.1613/jair.1.11640>.
- Peter H Schönemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1):1–10, 1966.
- Samuel L. Smith, David H. P. Turban, Steven Hamblin, and Nils Y. Hammerla. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. *CoRR*, abs/1702.03859, 2017. URL <http://arxiv.org/abs/1702.03859>.
- Samuel L. Smith, David H. P. Turban, Steven Hamblin, and Nils Y. Hammerla. Offline bilingual word vectors, orthogonal transformations and the inverted softmax. In *ICLR (Poster)*, 2017.
- Ivan Vulić and Anna Korhonen. On the role of seed lexicons in learning bilingual word embeddings. In *ACL*, 2016.
- Ivan Vulić and Marie-Francine Moens. Monolingual and cross-lingual information retrieval models based on (bilingual) word embeddings. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pp. 363–372, 2015.
- Ivan Vulić, Goran Glavaš, Roi Reichart, and Anna Korhonen. Do We Really Need Fully Unsupervised Cross-Lingual Embeddings? *arXiv e-prints*, art. arXiv:1909.01638, September 2019.
- Zihao Wang, Datong P. Zhou, Yong Zhang, Hao Wu, and Chenglong Bao. Wasserstein-fisher-rao document distance. *ArXiv*, abs/1904.10294, 2019.
- Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *Bernoulli*, 25:2620–2648, 2019.
- Chao Xing, Dong Wang, Chao Liu, and Yiye Lin. 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*, pp. 1006–1011, 2015.



Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. Earth mover's distance minimization for unsupervised bilingual lexicon induction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1934–1945, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1207. URL <https://aclanthology.org/D17-1207>.