

Constrained Density Matching and Modeling for Cross-lingual Alignment of Contextualized Representations

Wei Zhao

*Heidelberg Institute for Theoretical Studies
Fachbereich Informatik, Technische Universität Darmstadt*

WEI.ZHAO@H-ITS.ORG

Steffen Eger

*Technische Fakultät, Universität Bielefeld
Fachbereich Informatik, Technische Universität Darmstadt*

STEFFEN.EGER@UNI-BIELEFELD.DE

Editors: Emtiyaz Khan and Mehmet Gönen

Abstract

Multilingual representations pre-trained with monolingual data exhibit considerably unequal task performances across languages. Previous studies address this challenge with resource-intensive contextualized alignment, which assumes the availability of large parallel data, thereby leaving under-represented language communities behind. In this work, we attribute the data hungriness of previous alignment techniques to two limitations: (i) the inability to sufficiently leverage data and (ii) these techniques are not trained properly. To address these issues, we introduce supervised and unsupervised density-based approaches named Real-NVP and GAN-Real-NVP, driven by Normalizing Flow, to perform alignment, both dissecting the alignment of multilingual subspaces into density matching and density modeling. We complement these approaches with our validation criteria in order to guide the training process. Our experiments encompass 16 alignments, including our approaches, evaluated across 6 language pairs, synthetic data and 5 NLP tasks. We demonstrate the effectiveness of our approaches in the scenarios of limited and no parallel data. First, our supervised approach trained on 20k parallel data (sentences) mostly surpasses Joint-Align and InfoXML trained on over 100k parallel sentences. Second, parallel data can be removed without sacrificing performance when integrating our unsupervised approach in our bootstrapping procedure, which is theoretically motivated to enforce equality of multilingual subspaces. Moreover, we demonstrate the advantages of validation criteria over validation data for guiding supervised training¹.

Keywords: Multilingual Embeddings; Cross-lingual Alignment

1. Introduction

Multilingual text encoders such as m-BERT (Devlin et al., 2019) and XLM-R (Conneau et al., 2020) have been profiled as the de facto solutions to modeling languages at scale. However, research showed that such encoders pre-trained with monolingual data have failed to align multilingual subspaces, and exhibit strong language bias, i.e., the quality of these encoders largely differs across languages, particularly for dissimilar and low-resource languages (Pires et al., 2019; Zhao et al., 2021).

For that reason, supervised alignment techniques emerged, aiming to rectify multilingual representations post-hoc with cross-lingual supervision (Cao et al., 2020; Zhao et al., 2020; Chi et al., 2021), but previous studies are limited in scope to high-resource languages requiring large-scale par-

1. Our code and models are available at <https://github.com/AIPHES/DensityAlign>

allel data. In contrast, unsupervised alignment removing the dependence on parallel data allows for unlimited use in all languages (Artetxe et al., 2017, 2018). However, previous studies of unsupervised alignment focusing on static embeddings have not touched on contextualized representations.

In this work, we address cross-lingual alignment for contextualized representations, termed *contextualized alignment*, particularly in the scenarios of limited and no parallel data. In **supervised** settings, we identify two limitations responsible for the ineffectiveness of previous resource-intensive contextualized alignments: (i) the inability to sufficiently leverage data, i.e., that these techniques do not target the modeling of data density, and (ii) that they are not properly trained due to a lack of validation criteria—recent techniques, such as Wu and Dredze (2020) and Cao et al. (2020), have been trained for several epochs without access to any criteria for model selection, coming at the risk of being mistrained. To this end, we start by introducing a density-based, contextualized alignment, which dissects the alignment of multilingual subspaces into two sub-problems with one solution: density modeling and density matching, addressed by Normalizing Flows (Dinh et al., 2017). Second, in order to guide the training process, we present two validation criteria for model selection during training, and demonstrate the superiority of these criteria over validation data.

In **unsupervised settings**, aiming for unsupervised, contextualized alignment, we carry out density modeling and density matching in the form of adversarial learning (Goodfellow et al., 2014), and complement this learning process with the validation criteria mentioned previously to guide unsupervised training. Further, we identify a statistical issue of density matching in the unsupervised case: density matching only leads to a weak notion of equality of multilingual subspaces, *viz.*, equality in distribution. Accordingly, we present a bootstrapping procedure enhancing unsupervised alignment by promoting equality of multilingual subspaces. We evaluate our approaches across 6 language pairs, synthetic data and 5 NLP tasks. Our major findings are summarized as follows:

- With 20k parallel data we provided, our supervised alignment mostly surpasses Joint-Align (Cao et al., 2020) and InfoXLM (Chi et al., 2021) trained on much larger parallel data. This confirms the effectiveness of the conflation of density matching and density modeling as our alignment does. Second, our unsupervised alignment integrated in bootstrapping procedure rivals supervised counterparts, showing that parallel data can be removed without sacrificing performance. But we admit that these alignments, be it supervised or not, are poor in generalization (see §4.3), calling for an improvement in future work.
- Not only are validation criteria crucial for guiding unsupervised training, but also for supervised training. Given the performance on validation data and external tasks often correlates weakly, validation data is inappropriate for guiding supervised training. Above all, guiding contextualized alignment with validation criteria is challenging, as the model performances across tasks exhibit negative correlations in about 30% setups in our experiments, i.e., the better the alignment performs in one task, the worse it performs in the other. Thus, we base the evaluation of validation criteria on the model performances in all tasks. We find that validation criteria correlate much better than validation data (treated as criterion) with model performance on average across tasks for guiding supervised training.

2. Related Work

Recent advances in multilingual representations, such as m-BERT and XLM-R, boost the performance of cross-lingual NLP systems. However, such systems exhibit weak(er) performance for

dissimilar languages (Pires et al., 2019) and low-resource languages (Zhao et al., 2021). Accordingly, contextualized alignment emerged. Aldarmaki and Diab (2019) show that language-dependent rotation can linearly rectify m-BERT representations. Cao et al. (2020) find that jointly aligning multiple languages performs better. Zhao et al. (2020) show that removing language bias in multilingual representations mitigates the vector space misalignment between languages. More recently, Mengzhou et al. (2021) show that gradient-based alignment is effective for the languages not covered during pre-training in XLM-R. Alqahtani et al. (2021) use optimal transport to finetune multilingual representations, while Chi et al. (2021) finetune them with translation language modeling as the learning objective. However, these studies focused on supervised, resource-intensive alignment techniques and required from 250k to ca. 2M parallel sentences (or a large-scale analogy corpus) in each language pair for substantial improvement. As early attempts to remove the use for parallel data, Libovický et al. (2020) and Zhao et al. (2021) find applying vector space normalization is helpful to yield language-neutral representations. However, there lacks a thorough study on unsupervised, contextualized alignment for multilingual representations.

As for unsupervised alignment, previous studies have predominantly focused on static embeddings, which mostly rely on iterative procedures in two steps, aiming to derive bilingual lexicons as cross-lingual supervision: (i) inducing seed dictionaries with different approaches, such as adversarial learning (Lample et al., 2018), similarity based heuristics (Artetxe et al., 2018) and identical strings (Artetxe et al., 2017), and (ii) applying Procrustes to augment induced lexicons (Lample et al., 2018) in an iterative fashion.

In this work, we present a principled, iterative procedure to enhance our unsupervised alignment on contextualized representations, which employs density-based approaches to induce bilingual lexicons, and then applies our bootstrapping procedure, theoretically grounded in statistics for equality of multilingual subspaces, to iteratively augment lexicons. Lastly, we complement the iterative procedure with validation criteria to guide unsupervised training. We contrast our approaches with other domain adaptation techniques in Section 6.

3. Contextualized Alignment

Let two random variables X and Y with densities P_X and P_Y describe two populations of contextual word embeddings pertaining to two languages ℓ_1 and ℓ_2 , with Ω_{ℓ_1} and Ω_{ℓ_2} as two lexicons. Each occurrence of a word is associated to a separate entry in the lexicons. X maps all entries in Ω_{ℓ_1} to real-valued m -dimensional embedding vectors, denoted by $X : \Omega_{\ell_1} \rightarrow \mathbb{R}^m$, and similarly for Y . A bilingual lexicon Ω describes a set of translations between Ω_{ℓ_1} and Ω_{ℓ_2} .

Empirical inference. Assume a function $f : \mathbb{R}^m \rightarrow \mathbb{R}^m$ perfectly maps m -dimensional embedding vectors from X to Y . As standard in machine learning, a mapping function f_θ can be empirically inferred from data, with θ as model parameters. To this end, we assume data $\mathbf{M}_{\ell_1} \in \mathbb{R}^{n \times m}$ and $\mathbf{M}_{\ell_2} \in \mathbb{R}^{n \times m}$ are given, corresponding to two sets of contextual word embeddings with a common size of n for simplicity. Let a permutation matrix $\mathbf{P} \in \{0, 1\}^{n \times n}$ ($\mathbf{P}\mathbf{1}_n = \mathbf{1}_n$ and $\mathbf{P}^\top \mathbf{1}_n = \mathbf{1}_n$) be a realization of Ω , serving as cross-lingual supervision when available. A random variable \tilde{Y} with density $P_{\tilde{Y}}$ is a prediction of Y given X , i.e., $\tilde{Y} = f_{X \rightarrow Y}(X)$ where the subscript denotes mapping direction.

3.1. Supervised Alignment

When parallel data is available, a permutation matrix \mathbf{P} can be effortlessly induced from parallel data with word alignment tools (Dyer et al., 2013; Jalili Sabet et al., 2020). We introduce a density-based mapping function focusing on two components: density matching and density modeling. To do so, we start by depicting the alignment of multilingual subspaces in the form of density matching between $P_{\tilde{Y}}$ and P_Y :

$$\begin{aligned} \text{KL}(P_{\tilde{Y}}, P_Y) &= \text{CE}(P_{\tilde{Y}}, P_Y) - \mathbb{E}_{y \sim P_Y} [\log P_Y(y)] \\ &= \|f_{X \rightarrow Y}(\mathbf{M}_{\ell_1}) - \mathbf{P}\mathbf{M}_{\ell_2}\|^2 - \mathbb{E}_{y \sim P_Y} [\log P_X(f_{X \rightarrow Y}^{-1}(y)) |\det(\nabla_{\theta} f_{X \rightarrow Y}^{-1}(y))|] \end{aligned} \quad (1)$$

where $f_{X \rightarrow Y}$ is the trainable mapping function from X to Y . Given P_Y intractable to compute, previous supervised alignments always minimize the cross-entropy term alone by solving the least squares problem. Note that the density P_Y can be rewritten to $P_X(f_{X \rightarrow \tilde{Y}}^{-1}(y)) |\det(\nabla_{\theta} f_{X \rightarrow \tilde{Y}}^{-1}(y))|$ by using the change-of-variable rule (assuming f is an invertible function). However, the density P_Y is still intractable given the unknown density P_X . We overcome this by introducing a generative model named Real-NVP (Dinh et al., 2017) as use case of Normalizing Flows (Rezende and Mohamed, 2015). Real-NVP is a popular example of invertible neural networks, which can be thought of as a bijective function between two domains of data points (e.g., random noise and real data). Here, we use Real-NVP to address density estimation, i.e., inferring the unknown distribution of word embeddings X and Y from a normal distribution of random noise via the change-of-variable rule.

To do so, we introduce a latent variable $Z \sim \mathcal{N}(0, \mathbf{I})$ with the normal density P_Z to describe random noise. We then use Real-NVP to infer P_Y from P_Z , denoted by $P_Y(y) = P_Z(f_{Z \rightarrow Y}^{-1}(y)) |\det(\nabla_{\theta} f_{Z \rightarrow Y}^{-1}(y))|$ with $f_{Z \rightarrow Y}$ as a trainable mapping function from Z to Y . Lastly, we rewrite the entropy term in Eq. 1 to:

$$\mathbb{E}_{y \sim P_Y} [\log P_Y(y)] = \mathbb{E}_{y \sim P_Y} [\log \mathcal{N}(f_{Z \rightarrow Y}^{-1}(y), 0, \mathbf{I})] + \mathbb{E}_{y \sim P_Y} [\log |\det(\nabla_{\theta} f_{Z \rightarrow Y}^{-1}(y))|] \quad (2)$$

To consider density estimation (modeling) on both P_X and P_Y , we perform a dual form of density matching based on JS divergence. We omit the definition for simplicity. In §4, we refer to the above described approach as Real-NVP.

3.2. Unsupervised Alignment

When \mathbf{P} is not given due to a lack of parallel data, we apply adversarial learning to align the two densities $P_{\tilde{Y}}$ and P_Y . As standard in adversarial training, we involve a min-max game between two components to perform density matching: (a) a discriminator distinguishing source and target word embeddings after mapping them and (b) a mapping function aligning source and target word embeddings in order to fool the discriminator. We use a popular adversarial approach, the Wasserstein GAN (Arjovsky et al., 2017), which aligns the densities $P_{\tilde{Y}}$ and P_Y by minimizing the Earth Mover distance (EMD) between these densities. To better leverage data, we include density estimation (modeling) based on Real-NVP in the procedure of adversarial training, which maximizes the data likelihood of X and Y . Taken together, we denote our density-based learning objective by:

$$\text{EMD}(P_{\tilde{Y}}, P_Y) = \min_{f_{X \rightarrow Y}} \max_{h_{\phi}} \mathbb{E}_{y \sim P_Y} [h_{\phi}(y)] - \mathbb{E}_{\tilde{y} \sim P_{\tilde{Y}}} [h_{\phi}(\tilde{y})] + \mathbb{E}_{y \sim P_Y} [\log P_Y(y)] \quad (3)$$

where h_ϕ is a 1-Lipschitz constrained discriminator, $\tilde{y} = f_{X \rightarrow Y}(x)$ mapping X to Y . Note that $f_{X \rightarrow Y}$ is the composition of $f_{X \rightarrow Z}$ and $f_{Z \rightarrow Y}$, and the last entropy term aims to maximize the data log-likelihood of Y . As in the supervised case, we use a dual form of Eq. 3. In §4, we refer to the above described approach as GAN-Real-NVP.

Bootstrapping procedure. After adversarial learning Y and \tilde{Y} are ideally equal in distribution, denoted by $Y \stackrel{\text{dist}}{=} \tilde{Y}$. However, this is not sufficient. For instance, let $Y \sim \text{Uniform}(-1, 1)$ and $\tilde{Y} = -Y$. Clearly, Y and \tilde{Y} are equal in distribution, but they are identical only at the origin. Here, we derive the two following conditions that promote the equality of Y and \tilde{Y} and enhance unsupervised alignment.

Proposition 1 *Given $Y \stackrel{\text{dist}}{=} \tilde{Y}$, \tilde{Y} and Y are equal if one of the following conditions is met:*

- (i) $Y = \mathbf{U}\tilde{Y}$, where \mathbf{U} is invertible and $\mathbf{U}_{ij} \geq 0 \forall i, j$.
- (ii) $\text{cor}(\tilde{Y}_i, Y_i) = 1$ for $\forall i$, where \tilde{Y}_i and Y_i represent the i -th component in \tilde{Y} and Y .

Proof

(i) $P(Y \leq y) = P(\tilde{Y} \leq y)$ for all y , due to $Y \stackrel{\text{dist}}{=} \tilde{Y}$. If $Y = \mathbf{U}\tilde{Y}$, then $P(\tilde{Y} \leq y) = P(Y \leq y) = P(\mathbf{U}\tilde{Y} \leq y)$. If $\mathbf{U} \geq 0$, then $P(\mathbf{U}\tilde{Y} \leq y) = P(\tilde{Y} \leq \mathbf{U}^{-1}y)$. Thus, $P(\tilde{Y} \leq \mathbf{U}^{-1}y) = P(\tilde{Y} \leq y)$ for all y . This implies that $\mathbf{U} = \mathbf{I}$. Thus, $\tilde{Y} = Y$.

(ii) If $\text{cor}(\tilde{Y}_i, Y_i) = 1$ for $\forall i$, then $\text{Var}[(\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} - \frac{Y_i}{\sigma_{Y_i}})] = 0$, thus $\mathbb{E}[(\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} - \frac{Y_i}{\sigma_{Y_i}})^2] - \mathbb{E}[(\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} - \frac{Y_i}{\sigma_{Y_i}})]^2 = 0$. However, the second term equals to 0 by using $\mathbb{E}[(\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} - \frac{Y_i}{\sigma_{Y_i}})] = \frac{\mathbb{E}[\tilde{Y}_i]}{\sigma_{\tilde{Y}_i}} - \frac{\mathbb{E}[Y_i]}{\sigma_{Y_i}} = 0$ due to $\tilde{Y}_i \stackrel{\text{dist}}{=} Y_i$. Thus, $\mathbb{E}[(\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} - \frac{Y_i}{\sigma_{Y_i}})^2] = 0$, and this implies $\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} = \frac{Y_i}{\sigma_{Y_i}}$ since the non-negative $(\frac{\tilde{Y}_i}{\sigma_{\tilde{Y}_i}} - \frac{Y_i}{\sigma_{Y_i}})^2$ must be zero if its expectation is 0. Note that $\sigma_{\tilde{Y}_i} = \sigma_{Y_i}$ due to $\tilde{Y}_i \stackrel{\text{dist}}{=} Y_i$. This implies that $\tilde{Y}_i = Y_i$ for $\forall i$. ■

To design computational approaches meeting the above conditions, we introduce additional notation and the following lemma. Let $\mathbf{M}_X, \mathbf{M}_Y$ be embeddings from X and Y , and $\mathbf{M}_{\tilde{Y}} = f_\theta(\mathbf{M}_X)$.

Lemma 2 *If $\mathbf{M}_{\tilde{Y}}\mathbf{M}_{\tilde{Y}}^\top = \mathbf{M}_Y\mathbf{M}_Y^\top$ and $\mathbf{M}_{\tilde{Y}}$ is invertible, then $Y = \mathbf{U}\tilde{Y}$.*

Proof If $\mathbf{M}_{\tilde{Y}}\mathbf{M}_{\tilde{Y}}^\top = \mathbf{M}_Y\mathbf{M}_Y^\top$ and $\mathbf{M}_{\tilde{Y}}$ is invertible, then $\mathbf{M}_Y = \mathbf{M}_{\tilde{Y}}\mathbf{M}_{\tilde{Y}}^{-1}\mathbf{M}_Y$. Let $\mathbf{U} = \mathbf{M}_{\tilde{Y}}^{-1}\mathbf{M}_Y$. Then, $\mathbf{M}_Y = \mathbf{M}_{\tilde{Y}}\mathbf{U}$. If this holds for all $\mathbf{M}_{\tilde{Y}}$ and \mathbf{M}_Y , then $Y = \mathbf{U}\tilde{Y}$. ■

In the following, we describe our computational approaches, and then include them as **constraints** in the adversarial training in order to promote the equality of Y and \tilde{Y} . Lastly, we discuss the connection of these constraints with canonical correlation and language isomorphism.

Graph structure. We depict $\mathbf{M}_{\tilde{Y}}$ and \mathbf{M}_Y as m -dimensional vertices in two graphs, with $\mathbf{M}_{\tilde{Y}}\mathbf{M}_{\tilde{Y}}^\top$ and $\mathbf{M}_Y\mathbf{M}_Y^\top$ as the weighted adjacency matrices on these graphs. As Lemma 2 states, minimizing the difference between these adjacency matrices allows to meet Prop.1(i). Thus, the objective becomes:

$$\text{EMD}(P_{\tilde{Y}}, P_Y) + \|\mathbf{M}_{\tilde{Y}}\mathbf{M}_{\tilde{Y}}^\top - \mathbf{M}_Y\mathbf{M}_Y^\top\|^2 \quad (4)$$

However, we admit that Prop.1(i) cannot be strictly met, as guaranteeing $\mathbf{U} \geq 0$, i.e., $\mathbf{M}_{\tilde{Y}}^{-1}\mathbf{M}_Y \geq 0$ is not trivial. This might explain why *graph structure* is worse than *cross-correlation* in our experiments.

Cross-correlation. We maximize Pearson cross-correlation between de-meaned $\mathbf{M}_{\bar{Y}}$ and \mathbf{M}_Y in order to realize Prop.1(ii). The objective becomes:

$$\text{EMD}(P_{\bar{Y}}, P_Y) + \left\| \frac{\text{diag}(\mathbf{M}_{\bar{Y}}^T \mathbf{M}_Y)}{\text{diag}(\mathbf{M}_{\bar{Y}}^T \mathbf{M}_{\bar{Y}}) \text{diag}(\mathbf{M}_Y^T \mathbf{M}_Y)} - \vec{1} \right\|^2 \quad (5)$$

Concerning the construction of \mathbf{M}_Y and $\mathbf{M}_{\bar{Y}}$, we use CSLS (Lample et al., 2018) to induce them from monolingual data, and then update them in an iterative fashion with Algorithm 1.

Algorithm 1: Bootstrapping Procedure

Input: $\mathbf{M}_X, \mathbf{M}_Y \leftarrow$ population word embeddings of X and Y

Input: $n \leftarrow$ number of bootstrapping iterations \triangleright simulation: $n = 10$; real data: $n = 3$

Input: $f_{X \rightarrow Y} \leftarrow$ an identity matrix as initial mapping function

for $i \leftarrow 1$ **to** n **do**

$\mathbf{M}_{\bar{Y}} \leftarrow f_{X \rightarrow Y}(\mathbf{M}_X)$
 $\mathbf{P} \leftarrow \text{CSLS}(\mathbf{M}_Y, \mathbf{M}_{\bar{Y}})$ \triangleright induce permutation matrix
 $f_{X \rightarrow Y} \leftarrow \text{EMD}(P_{\bar{Y}}, P_Y) + g(\mathbf{M}_Y, \mathbf{P}\mathbf{M}_{\bar{Y}})$ \triangleright g is a function of our constraints (see Eq. 4+5)

end

Output: $f_{X \rightarrow Y}$

Connection with canonical correlation. Often, cross-correlation between random vectors are computed using Canonical Correlation Analysis (CCA). Research showed that CCA is useful to improve static embeddings, but it requires finding k primary canonical variables (Faruqui and Dyer, 2014). In contrast, our solution is much cheaper to compute cross-correlation without the need for canonical variables (see Eq. 5).

Connection with language isomorphism. In graph theory, two graphs are called isomorphic when the two corresponding adjacency matrices are permutation similar. According to Eq. 4, our solution aims to minimize the difference between adjacency matrices, and as such lays the foundation of graph isomorphism—which is termed language isomorphism in the multilingual community. Taken together, our solution allows for yielding isomorphic multilingual subspaces for non-isomorphic languages such as typologically dissimilar languages.

4. Experiments

4.1. Baselines and Our Approach

Supervised alignments. (a) Rotation (Aldarmaki and Diab, 2019; Zhao et al., 2021): a linear orthogonal-constrained transformation; (b) GBDD (Zhao et al., 2020): subtracting a global language bias vector from multilingual representations; (c) FCNN: an architecture that contains three fully-connected layers followed by a tanh activation function each; (d) Joint-Align (Cao et al., 2020): jointly aligning many languages via fine-tuning; (e) InfoXLM (Chi et al., 2021): finetuning multilingual representations with translation language modeling and contrastive learning; (f) our Real-NVP.

Unsupervised alignments. (a) MUSE: the unsupervised variant of Rotation (Lample et al., 2018); (b) VecMap: a heuristic unsupervised approach based on the assumption that word translations have

similar distributions on word similarities (Artetxe et al., 2018); (c) vector space normalization (Zhao et al., 2021): removing language-specific means and variances of multilingual representations. MUSE and VecMap are popular unsupervised alignments on static embeddings; (d) our GAN-Real-NVP. For bootstrapping procedure, we use the notation: [Method]+[Constraint], where [Method] is MUSE or GAN-Real-NVP, and [Constraint] is Cross-Correlation or Graph-Structure or Procrustes—known to enhance unsupervised alignment on static embeddings.

Except for Joint-Align, InfoXLM, and Normalization, the others are trained individually across language pairs.

4.2. Validation Criterion

We present two validation criteria, and compare them with no-criteria (i.e., training for several epochs) in both supervised and unsupervised settings. In particular, we induce the 30k most confident word translations from monolingual data with CSLS, and then compute the two following criteria on these word translations.

- *Semantic criterion* was proposed for guiding the training of unsupervised alignment on static embeddings. Lample et al. (2018) assemble the 10k most frequent source words and generate target translations of these words. Next, they average cosine similarities on these translation pairs treated as validation criterion.
- *Structural Criterion*: we compute the difference between two ordered lists of singular values obtained from source and target word embeddings pertaining to the 30k most confident word translations. This criterion was initially proposed to measure language isomorphism (Dubossarsky et al., 2020).

4.3. Simulation

Bilingual Lexicon Induction (BLI) is a popular internal task known to evaluate alignment on static embeddings, as it covers ca. 100 language pairs and focuses on the understanding of the alignment itself other than its impact on external tasks. In particular, BLI bases the induction of bilingual lexicons on static word embeddings, and compares the induced lexicons with gold lexicons.

However, contextual embeddings lack such evaluation tasks. As Artetxe et al. (2020) state, when not evaluated under similar conditions, the lessons learned from static embeddings cannot transfer to contextual ones. To this end, we perform simulation to construct synthetic data as the contextual extension of BLI (CBLI), which focuses on evaluating the alignment of multilingual subspaces of contextualized embeddings.

We split CBLI data to train, validation and test sets, and report Precision@K, as in BLI evaluation. Our creation procedure is two-fold: First, we sample source embeddings from a two-dimensional Gaussian (normal) mixture distribution, and then perform different transformations on them to produce target embeddings. By doing so, we mimic typologically dissimilar languages—see Figure 1.

As for the construction of simulation setups, we adjust three parameters: (a) the occurrence for a word k , $k \in \{5, 10, \dots, 100\}$ —we use 20 words in all setups; (b) the degree of language isomorphism $t \in \{1, \dots, 10\}$ —which mimics different language pairs and (c) the distance ϵ between embeddings in train and test sets, $\epsilon \in \{0, 0.2, \dots, 5\}$ —which reflects different similarities between train and test domains. For the i -th word, in order to reflect word occurrence, we sample contextualized embeddings from a normal distribution $\mathcal{N}(\mu_i, \mathbf{I})$ for train sets, and from $\mathcal{N}(\mu_i + \epsilon, \mathbf{I})$ for validation




Figure 1: Eight figures are constructed in simulation. Each depicts two languages pertaining to two subspaces, colored in blue and red. Each subspace consists of up to 3 densities with each representing a word. Each density contains a number of data points sampled from a two-dimensional Gaussian distribution, as a reflection of word occurrence.

and test sets, based on the insights from the visualized m-BERT space: different instances of a word appear to follow a normal distribution (Cao et al., 2020). μ_i denotes a mean vector sampled uniformly from $[-5, 5]$ for each component. For isomorphic languages ($t = 1$), we transform source into target embeddings with a rotation matrix. For non-isomorphic languages ($t > 1$), we alternate rotation with translation t times, assuming the more often we alternate, the more dissimilar two languages become.

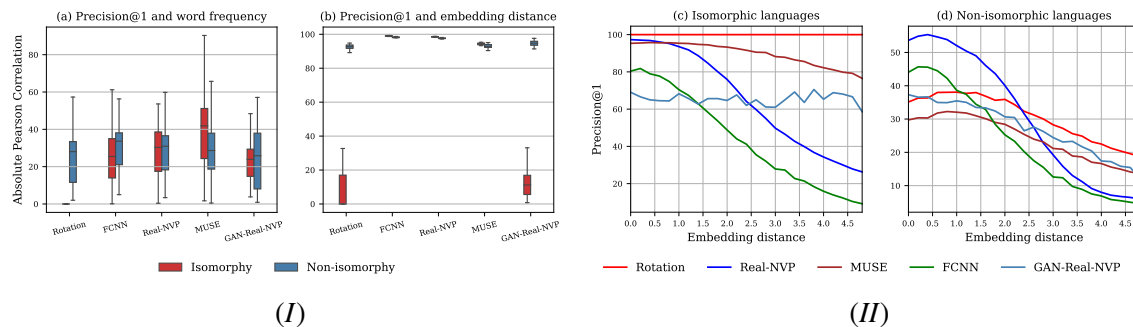


Figure 2: Absolute Pearson correlation between task performance and (a) word frequency (occurrence) and (b) similarity between train and test domains (the distance between embeddings on train and test sets). We set frequency bins $k \in \{5, 10, \dots, 100\}$, and similarity bins $\epsilon \in \{0, 0.2, \dots, 5\}$. We set t to 1 in all isomorphic settings, and t to 5 in non-isomorphic settings. (c)+(d) compares the generalization of approaches. Results are averaged across 10 runs.

Generalization to unseen words. Research showed that word frequency has a big impact on task performance for static embeddings (Czarnowska et al., 2019). However, Figure 2 (a)+(b) show that, in the contextual case, task performance often does not correlate with word frequency but strongly correlates with domain similarities between train and test sets. On a side note, Figure 2(b) shows that Rotation correlates poorly with embedding distance in “isomorphy”, but rather highly in “non-isomorphy”. This is because isomorphic spaces can be perfectly aligned via Rotation, independent of

the degree of embedding distance. For non-isomorphic spaces, the bigger the embedding distance is, the worse Rotation performs, which results in a high absolute correlation.

Analyses by Glavaš et al. (2019) showed that linear alignments are much better than non-linear counterparts on static embeddings. In the following, we contrast linear with non-linear alignments on contextualized embeddings, aiming to understand in which cases one is superior to another.

In isomorphic settings, Figure 2 (c) shows that linear alignments, Rotation and MUSE, clearly win in both supervised and unsupervised settings. This means a simple, linear transformation is sufficient to align vector spaces for isomorphic languages. We mark this as a sanity test, as languages mostly are non-isomorphic (Søgaard et al., 2018).

In non-isomorphic settings, Figure 2 (d) shows that non-linear alignments, Real-NVP and GAN-Real-NVP, win by a large margin when train and test domains are similar. However, when train and test domains are dissimilar, linear alignments are indeed better. As such, non-linear alignments suffer from the issue of generalization.

Overall, we show that alignment on static and contextual embeddings yield different conclusions on *word frequency* and *the superiority of linear over non-linear alignments*. By contrasting them, we hope to provide better understanding on each.




Figure 3: (I) shows how well two languages are aligned according to a visual introspection (subspace overlaps) and Precision@1; (II) compares unsupervised approaches (MUSE and GAN-Real-NVP) with the supervised counterparts (Rotation and Real-NVP) in non-isomorphic settings ($t = 5$). We set the occurrence per word k to 100.

Importance of bootstrapping procedure for unsupervised alignment. Figure 3 (I) shows how well two languages are aligned when train and test domains are similar. In this context, Real-NVP and GAN-Real-NVP win, and the resulting vector spaces are better overlapped (aligned) than others. This confirms the effectiveness of our density-based approaches. However, we still see a big performance gap between supervised and unsupervised approaches, especially for Real-NVP (64.3) vs. GAN-Real-NVP (51.3), notwithstanding large overlap in subspaces and small differences in model architectures. This confirms that density matching alone is not sufficient. Figure 3 (II) shows that after bootstrapping GAN-Real-NVP rivals Real-NVP. We also see similar results by contrasting Rotation with MUSE. Cross-correlation helps best in all cases, while graph-structure and Procrustes yield less consistent gains across approaches.

Overall, these results show that bootstrapping procedure plays a vital role in order for unsupervised alignments to rival supervised counterparts.

4.4. Experiments on Real Data

XTREME (Hu et al., 2020) has recently become popular for evaluating multilingual representations. However, it does not address word-level alignment as CBLI and BLI do, but rather focus on how multilingual representations impact cross-lingual systems. In this work, we evaluate both internal and external strengths of alignment, i.e., the internal alignment results on CBLI, and the impact of alignment on external tasks: (i) Align, RFEval and Tatoeba that require no supervised classifiers, and (ii) XNLI that requires a supervised classifier. We outline these tasks in the following:

- CBLI is the contextualized extension of BLI. Both contain a bilingual lexicon per language pair, but CBLI marks each occurrence of a word as an entry in lexicon. For each language pair, we extract 10k word translations from parallel sentences using FastAlign (Dyer et al., 2013). We report Precision@1. Note that we provide two complementary CBLI data: one is gold but simulated, while the other is real but contains noises.
- Alignment (Align) is a bilingual word retrieval task. Each language pair contains gold standard 2.5k word translations annotated by human experts. We use SimAlign (Jalili Sabet et al., 2020) to retrieve word translations from parallel sentences based on contextualized word embeddings. We report F-score that combines precision and recall.
- Reference-free evaluation (RFEval) measures the Pearson correlation between human and automatic judgments of translation quality. We use XMover (Zhao et al., 2019, 2020) to yield automatic judgment, which compares system translation with source sentence based on contextualized word embeddings. We exclude the target-side language model from XMover. Each language pair contains 3k source sentences.
- Tatoeba is a bilingual sentence retrieval task taken from XTREME. Each language pair contains 1k sentence translations. Given a source sentence, we retrieve the nearest translation from a pool of candidates based on cosine similarities between sentence embeddings. We report Precision@1.
- XNLI is a cross-lingual transfer task taken from XTREME, which aims to infer the relationship between a sentence pair of premise and hypothesis. Often, XNLI is evaluated in a zero-shot transfer setup, which measures the transfer ability from source to target languages, with cross-lingual systems trained on source language only. We report accuracy.

Tatoeba, CBLI and RFEval consist of six languages: German, Czech, Latvian, Finnish, Russian and Turkish, paired to English. In this work, we train alignments for these language pairs. Align considers two language pairs: German/Czech-to-English, and XNLI considers three: English-to-German/Russian/Turkish, as the other languages are not available. We consider two choices of multilingual representations: m-BERT and XLM-R.

Setup. To contrast supervised with unsupervised approaches, we consider two data scenarios: (i) limited parallel data and (ii) no parallel data. In case (i), we sample 20k (compared to ca. 250k often used in previous studies) parallel sentences from News-Commentary (Tiedemann, 2012) for Russian/Turkish-to-English, and from EuroParl (Koehn, 2005) for other languages. We use FastAlign to induce word translations from parallel sentences. Building upon these translations, we construct a permutation matrix \mathbf{P} as cross-lingual supervision. In case (ii), we unpair the word translations

obtained from (i) by removing the use for the permutation matrix. As such, we compare supervised and unsupervised approaches under similar conditions, *viz.*, with similar scale of data.

How to select the best model. We compare two choices of model selection: (i) CBLI as validation data and (ii) our validation criteria.

Figure 4 (I) shows that the results on CBLI and on other tasks correlate poorly (even negatively) in both supervised and unsupervised settings. This means validation data is inappropriate for guiding both supervised and unsupervised training.




Figure 4: Pearson correlation between task performances (I) and between validation criteria and task performance given by Real-NVP (II). Results are averaged across languages and encoders. For each task, we collect model performances and criteria scores over 20 epochs.

Settings	Alignments	$[-1, 0]$	$(0, 0.4]$	$(0.4, 1]$
Supervised	FCNN-20k	50%	0%	50%
	Real-NVP-20k	33%	17%	50%
Unsupervised	MUSE-20k	17%	66%	17%
	GAN-Real-NVP-20k	17%	50%	33%

Table 1: Correlation statistics: the last three columns split the Pearson’s ρ range into three intervals. Each entry denotes the percent of task pairs in which the correlation between model performances is in one of the intervals. For instance, the performances across tasks exhibit negative correlations in 17%-50% task pairs. For each task, we collect the model performances over 20 epochs. Results are averaged across language pairs and encoders.

Table 1 reports correlation statistics across approaches, showing that the model performances across tasks exhibit negative correlations in about 30% setups, *i.e.*, the better the alignment performs in one task, the worse it performs in the other. This means that (i) guiding contextualized alignment with validation criteria is challenging, and (ii) the evaluation of validation criteria should consider the performances in all tasks. As a running example, we evaluate the two validation criteria and validation data (CBLI) treated as criterion, based on Real-NVP. Figure 4 (II) shows that both validation criteria correlate much better than validation data with the model performances across tasks. Further, *semantic criterion* wins by a large margin (0.86 versus 0.44 and -0.12). This means *semantic criterion* is the best option for guiding Real-NVP. We also see similar results on other approaches. Thus, we adopt *semantic criterion* to perform model selection in all setups.

Overall, these results show that (i) not only are validation criteria important for guiding unsupervised training, but also for guiding supervised training, and (ii) the evaluation of validation criteria should be based on the model performances in all tasks.

4.5. Results on Real Data

Table 2 contrasts unsupervised with supervised approaches. For ease of reading, we provide the average results across languages, and break them down into individual languages in Table 4 (appendix).

Alignments	m-BERT				XLM-R			
	RFEval	Tatoeba	CBLI	Align	RFEval	Tatoeba	CBLI	Align
Original	27.23	49.35	50.9	61.54	26.42	63.40	48.58	59.77
<i>Supervised mapping functions</i>								
Rotation-20k	38.73	55.28	58.45	62.83	34.67	68.60	53.67	60.85
FCNN-20k (semantic criterion)	42.72	61.18	55.30	61.50	36.67	80.20	50.88	59.87
FCNN-20k (5 epochs)	38.40	58.97	54.02	61.02	33.50	77.98	50.48	59.50
Real-NVP-20k (semantic criterion)	42.32	62.87	57.62	62.59	44.17	80.08	61.63	62.84
Real-NVP-20k (5 epochs)	40.12	60.70	58.52	61.80	42.24	78.75	61.76	61.20
GBDD-20k	28.77	52.28	51.42	61.71	27.13	68.85	48.42	59.81
Joint-Align-100k (3 epochs)	41.23	59.13	64.67	62.30	-	-	-	-
InfoXLM-42GB (150K training steps)	-	-	-	-	37.60	76.10	60.80	62.94
<i>Unsupervised mapping functions</i>								
Normalization	30.08	61.28	54.88	62.54	39.52	79.75	59.03	62.55
VecMap-20k (~500 epochs)	30.77	55.00	64.42	62.50	-	-	-	-
MUSE-20k (5 epochs)	29.20	50.20	52.30	61.56	25.21	63.42	50.20	60.20
MUSE-20k (semantic criterion)	31.23	51.42	52.48	61.64	27.55	65.72	50.00	60.01
+ Cross-Correlation	35.25	52.90	52.87	62.63	32.05	69.23	49.80	60.49
+ Graph Structure	33.22	51.65	53.10	62.17	29.48	68.33	50.18	60.46
+ Procrustes	36.85	54.13	55.22	62.71	33.37	68.82	50.37	60.59
GAN-Real-NVP-20k (5 epochs)	32.24	59.10	56.79	61.80	39.61	77.77	60.83	60.90
GAN-Real-NVP-20k (semantic criterion)	33.90	61.20	57.03	62.33	41.72	79.67	61.00	62.81
+ Cross-Correlation	35.33	62.32	58.00	62.70	42.60	80.50	61.23	63.15
+ Graph Structure	34.32	61.82	56.65	62.52	41.55	80.02	60.83	62.99
+ Procrustes	36.93	53.95	56.05	62.79	33.78	67.95	51.60	60.51

Table 2: Results are averaged across language pairs. We bold numbers that significantly outperform others according to paired t-test. Joint-Align uses 100k parallel data per language pair; others only use 20k data. InfoXLM uses 42GB parallel data in total. Rotation, GBDD and Normalization with closed-form solutions do not require validation criteria for model selection.

Supervised settings. FCNN and Real-NVP training for 5 epochs are worse than those training with *semantic criterion* in almost all setups. This demonstrates the importance of validation criteria. We find that, though Joint-Align training with 100k parallel data wins on internal CBLI, it is worse than Real-NVP training with (i) merely 20k data and (ii) *semantic criterion* in the external tasks. This means Joint-Align overfits CBLI. We see similar results by contrasting VecMap with GAN-Real-NVP.

We see that the gains from all alignments on Align are much smaller than on others. This might be because SimAlign (used to induce word alignments) or the dataset cannot recognize the improved contextualized embeddings after alignment.

Real-NVP seems the strongest approach, which helps considerably for both m-BERT and XLM-R and surpasses recent InfoXML by a large margin. InfoXML training with much larger parallel data for 150k training steps cannot show advantages in the absence of validation criteria. Note that we do not apply validation criteria to Joint-Align and InfoXML for further improvements, as these resource-intensive approaches have not been designed for low-resource languages, e.g., that the improvements by Joint-Align appear to vanish in the setup of limited parallel data [Cao et al. \(2020\)](#).

Unsupervised settings. Validation criteria are crucial: MUSE and GAN-Real-NVP with *semantic criterion* largely outperform those training for 5 epochs.

Much unlike the results in simulation, we see small performance gaps between supervised and unsupervised approaches, such as the gap between Real-NVP and GAN-Real-NVP (2 points vs 13 points in simulation). Thus, it is not surprising that the gains from the bootstrapping procedure are small in these tasks. Overall, we see *cross-correlation* is better than *graph structure* on MUSE and GAN-Real-NVP. The results for Procrustes are similar as in simulation—it improves MUSE but harms GAN-Real-NVP. GAN-Real-NVP training with *semantic criterion* and *cross-correlation* always wins, rivaling the best supervised approach Real-NVP.

Overall, these supervised and unsupervised results show that (i) validation criterion plays an essential role; (ii) density-based approaches targeting density matching and density modeling are effective in both supervised and unsupervised settings, and (iii) after bootstrapping unsupervised approaches are able to rival supervised counterparts.

Alignments	m-BERT		
	DE	RU	TR
Original	70.3	68.2	60.0
Rotation-20k	70.6	68.1	60.5
FCNN-20k (semantic criterion)	70.8	68.1	59.9
Real-NVP-20k (semantic criterion)	73.6	72.7	62.9
Joint-Align-100k (3 epochs)	72.9	72.1	62.4
GAN-R-NVP-20k (semantic criterion)	73.5	72.2	61.5
+ Cross-Correlation	73.4	72.1	61.3
+ Graph Structure	73.4	72.3	61.2
+ Procrustes	70.4	68.3	60.2

Table 3: Results on XNLI in a zero-shot cross-lingual setup from English to German/Russian/Turkish. After rectifying m-BERT with alignments we finetune m-BERT (coupled with a supervised classifier) on XNLI English train data. We restrict the evaluation to 3 languages, as the other languages on XNLI are not covered in our alignments.

Downstream Task. Table 3 shows the impacts of our approaches coupled with a supervised classifier to perform zero-shot text classification on the downstream task XNLI. While Rotation and FCNN yield better results in Table 2, their impacts vanish on XNLI. This might be because (i) rectifying m-BERT with 20k parallel data is not adequate to reflect improvements on downstream tasks, or (ii) alignment results may be orthogonal to downstream zero-shot performance. However, Real-NVP and GAN-Real-NVP trained on the same scale of data with Rotation and FCNN exhibit strong impacts on XNLI, on par with Joint-Align trained with 100k parallel data. Thus, 20k data is sufficient for our approaches to yield improvements on XNLI.

5. Conclusion

Given resource and typology disparities across languages, multilingual representations exhibit unequal capabilities between languages. Research showed that contextualized alignments overcome this challenge by producing language-agnostic representations. However, these techniques demand large parallel data, and thus cannot address the data scarcity issue in low-resource languages. Our contributions in this work are manifold. We start by introducing supervised and unsupervised density-based approaches, Real-NVP and GAN-Real-NVP, both dissecting the alignment of multilingual subspaces into density matching and density modeling in order to sufficiently leverage data. Second, we investigate the usefulness of validation criteria for guiding the training process of our approaches. Further, we present a bootstrapping procedure to enhance our unsupervised approach, which is theoretically grounded for promoting equality of multilingual subspaces. We demonstrated the effectiveness of our alignments in the scenarios of limited and no parallel data. With 20k parallel data we provided, our supervised approach mostly outperforms Joint-Align and InfoXLM trained on much larger parallel data. Next, we showed that validation criteria are imperative for guiding both supervised and unsupervised training. Finally, we demonstrated that parallel data could be removed without the loss of model performances after integrating our unsupervised approach in the bootstrapping procedure.

6. Broader Impact

As a class of domain adaptation techniques, density-based approaches have been shown useful in a range of cross-domain applications, such as image-captioning (Mahajan et al., 2020), image-to-image translation (Grover et al., 2020), alignment on static embeddings (Zhou et al., 2019) and machine translation (Setiawan et al., 2020). In this work, we showed that (i) density-based approaches could overfit validation data in the absence of validation criteria, and are weak in generalization (see §4.3), but (ii) bootstrapping procedures can improve these density-based approaches. While our analyses are limited in scope to contextualized alignment as the only cross-domain application, we hope that our results fuel future research towards effective domain adaptation techniques in other applications.

Acknowledgments

We thank the anonymous reviewers for their thoughtful comments that greatly improved the final texts. We also thank Dan Liu for her early experiments in the word alignment task. This work has been supported by the German Research Foundation as part of the Research Training Group Adaptive Preparation of Information from Heterogeneous Sources (AIPHES) at the Technische Universität Darmstadt under grant No. GRK 1994/1 and the Klaus Tschira Foundation, Heidelberg, Germany.

References

- Hanan Aldarmaki and Mona Diab. Context-aware cross-lingual mapping. In *NAACL*, 2019.
- Sawsan Alqahtani, Garima Lalwani, Yi Zhang, Salvatore Romeo, and Saab Mansour. Using optimal transport as alignment objective for fine-tuning multilingual contextualized embeddings. In *EMNLP*, 2021.

- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *ICML*, 2017.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. Learning bilingual word embeddings with (almost) no bilingual data. In *ACL*, Vancouver, Canada, July 2017.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *ACL*, Melbourne, Australia, July 2018.
- Mikel Artetxe, Sebastian Ruder, Dani Yogatama, Gorka Labaka, and Eneko Agirre. A call for more rigor in unsupervised cross-lingual learning. In *ACL*, pages 7375–7388, Online, July 2020.
- Steven Cao, Nikita Kitaev, and Dan Klein. Multilingual alignment of contextual word representations. In *ICLR*, 2020.
- Zewen Chi, Li Dong, Furu Wei, Nan Yang, Saksham Singhal, Wenhui Wang, Xia Song, Xian-Ling Mao, Heyan Huang, and Ming Zhou. InfoXLM: An information-theoretic framework for cross-lingual language model pre-training. In *NAACL*, 2021.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised cross-lingual representation learning at scale. In *ACL*, Online, July 2020.
- Paula Czarrowska, Sebastian Ruder, Edouard Grave, Ryan Cotterell, and Ann Copestake. Don’t forget the long tail! a comprehensive analysis of morphological generalization in bilingual lexicon induction. In *EMNLP*, 2019.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. In *ICLR*, 2017.
- Haim Dubossarsky, Ivan Vulić, Roi Reichart, and Anna Korhonen. The secret is in the spectra: Predicting cross-lingual task performance with spectral similarity measures. In *EMNLP*, 2020.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. A simple, fast, and effective reparameterization of IBM model 2. In *NAACL*, 2013.
- Manaal Faruqui and Chris Dyer. Improving vector space word representations using multilingual correlation. In *EACL*, 2014.
- 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 *ACL*, Florence, Italy, July 2019.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- Aditya Grover, Christopher Chute, Rui Shu, Zhangjie Cao, and Stefano Ermon. Alignflow: Cycle consistent learning from multiple domains via normalizing flows. In *AAAI*, 2020.

- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. XTREME: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *ICML*, 2020.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. SimAlign: High quality word alignments without parallel training data using static and contextualized embeddings. In *Findings of EMNLP*, 2020.
- Philipp Koehn. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5. Citeseer, 2005.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. Word translation without parallel data. In *ICLR*, 2018.
- Jindřich Libovický, Rudolf Rosa, and Alexander Fraser. On the language neutrality of pre-trained multilingual representations. In *Findings of EMNLP*, Online, November 2020.
- Shweta Mahajan, Iryna Gurevych, and Stefan Roth. Latent normalizing flows for many-to-many cross-domain mappings. In *ICLR*, 2020.
- Xia Mengzhou, Guoqing Zheng, Subhabrata Mukherjee, Milad Shokouhi, Graham Neubig, and Ahmed Hassan Awadallah. Metaxl: Meta representation transformation for low-resource cross-lingual learning. 2021.
- Telmo Pires, Eva Schlinger, and Dan Garrette. How multilingual is multilingual BERT? In *ACL*, Florence, Italy, July 2019.
- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. In *ICML*, 2015.
- Hendra Setiawan, Matthias Sperber, Udhyakumar Nallasamy, and Matthias Paulik. Variational neural machine translation with normalizing flows. In *ACL*, July 2020.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. On the limitations of unsupervised bilingual dictionary induction. In *ACL*, 2018.
- Jörg Tiedemann. Parallel data, tools and interfaces in opus. In *LREC*, 2012.
- Shijie Wu and Mark Dredze. Do explicit alignments robustly improve multilingual encoders? In *EMNLP*, Online, November 2020.
- Wei Zhao, Maxime Peyrard, Fei Liu, Yang Gao, Christian M. Meyer, and Steffen Eger. MoverScore: Text generation evaluating with contextualized embeddings and earth mover distance. In *EMNLP*, Hong Kong, China, November 2019.
- Wei Zhao, Goran Glavaš, Maxime Peyrard, Yang Gao, Robert West, and Steffen Eger. On the limitations of cross-lingual encoders as exposed by reference-free machine translation evaluation. In *ACL*, 2020.
- Wei Zhao, Steffen Eger, Johannes Bjerva, and Isabelle Augenstein. Inducing language-agnostic multilingual representations. In **SEM*, Online, August 2021.
- Chunting Zhou, Xuezhe Ma, Di Wang, and Graham Neubig. Density matching for bilingual word embedding. In *NAACL*, Minneapolis, Minnesota, June 2019.