

Cross-language Sentence Selection via Data Augmentation and Rationale Training

Yanda Chen¹, Chris Kedzie¹, Suraj Nair², Petra Galuščáková², Rui Zhang³,
Douglas W. Oard², Kathleen McKeown¹

¹Columbia University, ²University of Maryland, ³Penn State University

yc3384@columbia.edu, {kedzie, kathy}@cs.columbia.edu
{srnair, petra, oard}@umd.edu, rmz5227@psu.edu

Abstract

This paper proposes an approach to cross-language sentence selection in a low-resource setting. It uses data augmentation and negative sampling techniques on noisy parallel sentence data to directly learn a cross-lingual embedding-based query relevance model. Results show that this approach performs as well as or better than multiple state-of-the-art machine translation + monolingual retrieval systems trained on the same parallel data. Moreover, when a rationale training secondary objective is applied to encourage the model to match word alignment hints from a phrase-based statistical machine translation model, consistent improvements are seen across three language pairs (English-Somali, English-Swahili and English-Tagalog) over a variety of state-of-the-art baselines.

1 Introduction

Sentence-level query relevance prediction is important for downstream tasks such as query-focused summarization and open-domain question answering; accurately pinpointing sentences containing information that is relevant to the query is critical to generating a responsive summary/answer (e.g., Baumel et al. (2016, 2018)). In this work, we focus on sentence-level query relevance prediction in a cross-lingual setting, where the query and sentence collection are in different languages and the sentence collection is drawn from a low-resource language. Our approach enables English speakers (e.g., journalists) to find relevant information expressed in local sources (e.g., local reaction to the pandemic and vaccines in Somalia).

While we can use machine translation (MT) to translate either the query or each sentence into a common language, and then use a monolingual Information Retrieval (IR) system to find relevant sentences, work on Probabilistic Structured Queries

(PSQ) (Darwish and Oard, 2003) has shown that the performance of such MT+IR pipelines is hindered by errors in MT. As is well known, complete translation of the sentence collection is not necessary. Inspired by previous work (Vulić and Moens, 2015), we go a step further and propose a simple cross-lingual embedding-based model that avoids translation entirely and directly predicts the relevance of a query-sentence pair (where the query and sentence are in different languages).

For training, we treat a sentence as *relevant* to a query if there exists a translation equivalent of the query in the sentence. Our definition of relevance is most similar to the lexical-based relevance used in Gupta et al. (2007) and Baumel et al. (2018) but our query and sentence are from different languages. We frame the task as a problem of finding sentences that are relevant to an input query, and thus, we need relevance judgments for query-sentence pairs. Our focus, however, is on low-resource languages where we have no sentence-level relevance judgments with which to train our query-focused relevance model. We thus leverage noisy parallel sentence collections previously collected from the web. We use a simple data augmentation and negative sampling scheme to generate a labeled dataset of relevant and irrelevant pairs of queries and sentences from these noisy parallel corpora. With this synthetic training set in hand, we can learn a supervised cross-lingual embedding space.

While our approach is competitive with pipelines of MT-IR, it is still sensitive to noise in the parallel sentence data. We can mitigate the negative effects of this noise if we first train a phrase-based statistical MT (SMT) model on the same parallel sentence corpus and use the extracted word alignments as additional supervision. With these alignment hints, we demonstrate consistent and significant improvements over neural and statistical MT+IR (Niu et al., 2018; Koehn et al., 2007; Heafield, 2011),

three strong cross-lingual embedding-based models (Bivec (Luong et al., 2015), SID-SGNS (Levy et al., 2017), MUSE (Lample et al., 2018)), a probabilistic occurrence model (Xu and Weischedel, 2000), and a multilingual pretrained model XLM-RoBERTa (Conneau et al., 2020). We refer to this secondary training objective as *rationale training*, inspired by previous work in text classification that supervises attention over rationales for classification decisions (Jain and Wallace, 2019).

To summarize, our contributions are as follows. We (i) propose a data augmentation and negative sampling scheme to create a synthetic training set of cross-lingual query-sentence pairs with binary relevance judgements, and (ii) demonstrate the effectiveness of a Supervised Embedding-based Cross-Lingual Relevance (SECLR) model trained on this data for low-resource sentence selection tasks on text and speech. Additionally, (iii) we propose a rationale training secondary objective to further improve SECLR performance, which we call SECLR-RT. Finally, (iv) we conduct training data ablation and hubness studies that show our method’s applicability to even lower-resource settings and mitigation of hubness issues (Dinu and Baroni, 2015; Radovanović et al., 2010). These findings are validated by empirical results of experiments in a low-resource sentence selection task, with English queries over sentence collections of text and speech in Somali, Swahili, and Tagalog.

2 Related Work

Query-focused Sentence Selection Sentence-level query relevance prediction is important for various downstream NLP tasks such as query-focused summarization (Baumel et al., 2016, 2018; Feigenblat et al., 2017) and open-domain question answering (Chen et al., 2017; Dhingra et al., 2017; Kale et al., 2018). Such applications often depend on a sentence selection system to provide attention signals on which sentences to focus upon to generate a query-focused summary or answer a question.

Cross-language Sentence Selection A common approach to cross-language sentence selection is to use MT to first translate either the query or the sentence to the same language and then perform standard monolingual IR (Nie, 2010). The risk of this approach is that errors in translation cascade to the IR system.

As an alternative to generating full translations, PSQ (Darwish and Oard, 2003) uses word-

alignments from SMT to obtain weighted query term counts in the passage collection. In other work, Xu and Weischedel (2000) use a 2-state hidden Markov model (HMM) to estimate the probability that a passage is relevant given the query.

Cross-lingual Word Embeddings Cross-lingual embedding methods perform cross-lingual relevance prediction by representing query and passage terms of different languages in a shared semantic space (Vulić and Moens, 2015; Litschko et al., 2019, 2018; Joulin et al., 2018). Both supervised approaches trained on parallel sentence corpora (Levy et al., 2017; Luong et al., 2015) and unsupervised approaches with no parallel data (Lample et al., 2018; Artetxe et al., 2018) have been proposed to train cross-lingual word embeddings.

Our approach differs from previous cross-lingual word embedding methods in two aspects. First, the focus of previous work has mostly been on learning a distributional word representation where translation across languages is primarily shaped by syntactic or shallow semantic similarity; it has not been tuned specifically for cross-language sentence selection tasks, which is the focus of our work.

Second, in contrast to previous supervised approaches that train embeddings directly on a parallel corpus or bilingual dictionary, our approach trains embeddings on an *artificial labeled dataset* augmented from a parallel corpus and directly represents relevance across languages. Our data augmentation scheme to build a relevance model is inspired by Boschee et al. (2019), but we achieve significant performance improvement by incorporating rationale information into the embedding training process and provide detailed comparisons of performance with other sentence selection approaches.

Trained Rationale Previous research has shown that models trained on classification tasks sometimes do not use the correct *rationale* when making predictions, where a rationale is a mechanism of the classification model that is expected to correspond to human intuitions about salient features for the decision function (Jain and Wallace, 2019). Research has also shown that incorporating human rationales to guide a model’s attention distribution can potentially improve model performance on classification tasks (Bao et al., 2018). Trained rationales have also been used in neural MT (NMT); incorporat-

ing alignments from SMT to guide NMT attention yields improvements in translation accuracy (Chen et al., 2016).

3 Methods

We first describe our synthetic training set generation process, which converts a parallel sentence corpus for MT into cross-lingual query-sentence pairs with binary relevance judgements for training our SECLR model. Following that, we detail our SECLR model and finish with our method for rationale training with word alignments from SMT.

3.1 Training Set Generation Algorithm

Relevant query/sentence generation. Assume we have a parallel corpus of bilingual sentence pairs equivalent in meaning. Let (E, S) be one such sentence pair, where E is in the query language (in our case, English) and S is in the retrieval collection language (in our case, low-resource languages). For every unigram q in E that is not a stopword, we construct a positive relevant sample by viewing q as a query and S as a relevant sentence. Because sentences E and S are (approximately) equivalent in meaning, we know that there likely exists a translation equivalent of q in the sentence S and so we label the (q, S) pair as relevant (i.e. $r = 1$).

For example, one English-Somali sentence pair is E =“true president gaas attend meeting copenhagen”, S =“ma runbaa madaxweyne gaas baaqday shirka copenhagen” (stopwords removed). By extracting unigrams from E as queries, we generate the following positive examples: (q =“true”, S , $r = 1$), (q =“president”, S , $r = 1$), (q =“gaas”, S , $r = 1$), ..., (q =“copenhagen”, S , $r = 1$).

We generate the positive half of the training set by repeating the above process for every sentence pair in the parallel corpus. We limit model training to unigram queries since higher order ngrams appear fewer times and treating them independently reduces the risk of over-fitting. However, our model processes multi-word queries during evaluation, as described in Section 3.2.

Irrelevant query/sentence generation. Since learning with only positive examples is a challenging task, we opt to create negative examples, i.e. tuples $(q, S, r = 0)$, via negative sampling. For each positive sample $(q, S, r = 1)$, we randomly select another sentence pair (E', S') from the parallel corpus. We then check whether S' is relevant to q or not. Note that both the query q and sentence

E' are in the same language, so checking whether q or a synonym can be found in E' is a monolingual task. If we can verify that there is no direct match or synonym equivalent of q in E' then by transitivity it is unlikely there exists a translation equivalent in S' , making the pair (q, S') a negative example. To account for synonymy when we check for matches, we represent q and the words in E' with pretrained word embeddings. Let $w_q, w_{q'} \in \mathbb{R}^d$ be the embeddings associated with q and the words $q' \in E'$. We judge the pair (q, S') to be irrelevant (i.e. $r = 0$) if:

$$\max_{q' \in E'} \cos\text{-sim}(w_q, w_{q'}) \leq \lambda_1$$

where λ_1 is a parameter. We manually tuned the relevance threshold λ_1 on a small development set of query-sentence pairs randomly generated by the algorithm, and set $\lambda_1 = 0.4$ to achieve highest label accuracy on the development set. If (q, S') is not relevant we add $(q, S', r = 0)$ to our synthetic training set, otherwise we re-sample (E', S') until a negative sample is found. We generate one negative sample for each positive sample to create a balanced dataset.

For example, if we want to generate a negative example for the positive example (q =“meeting”, S =“ma runbaa madaxweyne gaas baaqday shirka copenhagen”, $r = 1$), we randomly select another sentence pair (E' =“many candidates competing elections one hopes winner”, S' =“musharraxiin tiro badan sidoo u tartamaysa doorashada wuxuuna mid kasta rajo qabaa guusha inay dhinacisa ahaato”) from the parallel corpus. To check whether q =“meeting” is relevant to S' , by transitivity it suffices to check whether q =“meeting” or a synonym is present in E' , a simpler monolingual task. If q is irrelevant to S' , we add $(q, S', r = 0)$ as a negative example.

3.2 Cross-Lingual Relevance Model

We propose SECLR, a model that directly makes relevance classification judgments for queries and sentences of different languages without MT as an intermediate step by learning a cross-lingual embedding space between the two languages. Not only should translation of equivalent words in either language map to similar regions in the embedding space, but dot products between query and sentence words should be correlated with the probability of relevance. We assume the training set generation process (Section 3.1) provides us with a corpus of n query-sentence pairs along

with their corresponding relevance judgements, i.e. $\mathcal{D} = \{(q_i, S_i, r_i)\}_{i=1}^n$. We construct a bilingual vocabulary $\mathcal{V} = \mathcal{V}_Q \cup \mathcal{V}_S$ and associate with it a matrix $W \in \mathbb{R}^{d \times |\mathcal{V}|}$ where $w_x = W_{\cdot, x}$ is the word embedding associated with word $x \in \mathcal{V}$.

When the query is a unigram q (which is true by design in our training data \mathcal{D}), we model the probability of relevance to a sentence S as:

$$p(r = 1|q, S; W) = \sigma \left(\max_{s \in S} w_q^\top w_s \right)$$

where σ denotes the logistic sigmoid ($\sigma(x) = 1 / (1 + \exp(-x))$).

In our evaluation setting, the query is very often a phrase $Q = [q_1, \dots, q_{|Q|}]$. In this case, we require *all* query words to appear in a sentence in order for a sentence to be considered as relevant. Thus, we modify our relevance model to be:

$$p(r = 1|Q, S; W) = \sigma \left(\min_{q \in Q} \max_{s \in S} w_q^\top w_s \right)$$

Our only model parameter is the embedding matrix W which is initialized with pretrained monolingual word embeddings and learned via minimization of the cross entropy of the relevance classification task:

$$\mathcal{L}_{rel} = -\log p(r|q, S; W)$$

3.3 Guided Alignment with Rationale Training

We can improve SECLR by incorporating additional alignment information as a secondary training objective, yielding SECLR-RT. Our intuition is that after training, the word $\hat{s} = \arg \max_{s \in S} w_s^\top w_q$ should correspond to a translation of q . However, it is possible that \hat{s} simply co-occurs frequently with the true translation in our parallel data but its association is coincidental or irrelevant outside the training contexts. We use alignment information to correct for this. We run two SMT word alignment models, GIZA++ (Och and Ney, 2003) and Berkeley Aligner (Haghighi et al., 2009), on the original parallel sentence corpus. The two resulting alignments are concatenated as in Zbib et al. (2019) to estimate a unidirectional probabilistic word translation matrix $A \in [0, 1]^{|\mathcal{V}_Q| \times |\mathcal{V}_S|}$, such that A maps each word in the query language vocabulary to a list of document language words with different probabilities, i.e. $A_{q,s}$ is the probability of translating q to s and $\sum_{s \in \mathcal{V}_S} A_{q,s} = 1$.

For each relevant training sample, i.e. $(q, S, r = 1)$, we create a *rationale distribution* $\rho \in [0, 1]^{|S|}$

which is essentially a re-normalization of possible query translations found in S and represents our intuitions about which words $s \in S$ that q should be most similar to in embedding space, i.e.

$$\rho_s = \frac{A_{q,s}}{\sum_{s' \in S} A_{q,s'}}$$

for $s \in S$. We similarly create a distribution under our model, $\alpha \in [0, 1]^{|S|}$, where

$$\alpha_s = \frac{\exp(w_q^\top w_s)}{\sum_{s' \in S} \exp(w_q^\top w_{s'})}$$

for $s \in S$. To encourage α to match ρ , we impose a Kullback–Leibler (KL) divergence penalty, denoted as:

$$\mathcal{L}_{rat} = \text{KL}(\rho || \alpha)$$

to our overall loss function. The total loss for a single positive sample then will be a weighted sum of the relevance classification objective and the KL divergence penalty, i.e.

$$\mathcal{L} = \mathcal{L}_{rel} + \lambda_2 \mathcal{L}_{rat}$$

where λ_2 is a relative weight between the classification loss and rationale similarity loss.

Note that we do not consider rationale loss for the following three types of samples: negative samples, positive samples where the query word is not found in the translation matrix, and positive samples where none of the translations of the query in the matrix are present in the source sentence.

4 Experiments

4.1 Dataset Generation from Parallel Corpus

The parallel sentence data for training our proposed method and all baselines includes the parallel data provided in the BUILD collections of both the MATERIAL¹ and LORELEI (Christianson et al., 2018) programs for three low resource languages: Somali (SO), Swahili (SW), and Tagalog (TL) (each paired with English). Additionally, we include in our parallel corpus publicly available resources from OPUS (Tiedemann, 2012), and lexicons mined from Panlex (Kamholz et al., 2014) and Wiktionary.² Statistics of these parallel corpora and augmented data are shown in Table 1 and Table 2, respectively. Other preprocessing details are in Appendix A.

¹<https://www.iarpa.gov/index.php/research-programs/material>

²<https://dumps.wikimedia.org/>

	EN-SO	EN-SW	EN-TL
# sents.	69,818	251,928	232,166
EN tkn.	1,827,826	1,946,556	2,553,439
LR tkn.	1,804,428	1,848,184	2,682,076

Table 1: Parallel corpus statistics; “EN tkn.” refers to number of English tokens in the parallel corpus; “LR tkn.” refers to number of low-resource tokens (Somali, Swahili, Tagalog) in the parallel corpus.

Lang. Pair	Augmented Dataset Size
EN-SO	1,649,484
EN-SW	2,014,838
EN-TL	2,417,448

Table 2: Augmented dataset statistics; “augmented dataset size” refers to total number of positive and negative query-sentence samples in the augmented dataset.

4.2 Query Sets and Evaluation Sets

We evaluate our sentence-selection model on English (EN) queries over three collections in SO, SW, and TL recently made available as part of the IARPA MATERIAL program. In contrast to our training data which is synthetic, our evaluation datasets are human-annotated for relevance between real-world multi-domain queries and documents. For each language there are three partitions (Analysis, Dev, and Eval), with the former two being smaller collections intended for system development, and the latter being a larger evaluation corpus. In our main experiments we do not use Analysis or Dev for development and so we report results for all three (the ground truth relevance judgements for the TL Eval collection have not been released yet so we do not report Eval for TL). See Table 3 for evaluation statistics. All queries are text. The speech documents are first transcribed with an ASR system (Ragni and Gales, 2018), and the 1-best ASR output is used in the sentence selection task. Examples of the evaluation datasets are shown in Appendix B. We refer readers to Rubino (2020) for further details about MATERIAL test collections used in this work.

While our model and baselines work at the sentence-level, the MATERIAL relevance judgements are only at the document level. Following previous work on evaluation of passage retrieval, we aggregate our sentence-level relevance scores to obtain document-level scores (Kaszkiel and Zo-

bel, 1997; Wade and Allan, 2005; Fan et al., 2018; Inel et al., 2018; Akkalyoncu Yilmaz et al., 2019). Given a document $D = [S_1, \dots, S_{|D|}]$, which is a sequence of sentences, and a query Q , following Liu and Croft (2002) we assign a relevance score by:

$$\hat{r} = \max_{S \in D} p(r = 1 | Q, S; W)$$

4.3 Experiment Settings

We initialize English word embeddings with word2vec (Mikolov et al., 2013), and initialize SO/SW/TL word embeddings with FastText (Grave et al., 2018). For training we use a SparseAdam (Kingma and Ba, 2015) optimizer with learning rate 0.001. The hyperparameter λ_2 in Section 3.3 is set to be 3 so that \mathcal{L}_{rel} and $\lambda_2 \mathcal{L}_{rat}$ are approximately on the same scale during training. More details on experiments are included in Appendix C.

4.4 Baselines

Cross-Lingual Word Embeddings. We compare our model with three other cross-lingual embedding methods, Bivec (Luong et al., 2015), MUSE (Lample et al., 2018), and SID-SGNS (Levy et al., 2017). Bivec and SID-SGNS are trained using the same parallel sentence corpus as the dataset generation algorithm used to train SECLR; thus, Bivec and SID-SGNS are trained on parallel sentences while SECLR is trained on query-sentence pairs derived from that corpus. We train MUSE with the bilingual dictionary from Wiktionary that is used in previous work (Zhang et al., 2019). The SO-EN, SW-EN and TL-EN dictionaries have 7633, 5301, and 7088 words respectively. Given embeddings W' from any of these methods, we compute sentence level relevance scores similarly to our model but use the cosine similarity:

$$p(r = 1 | Q, S; W') = \min_{q \in Q} \max_{s \in S} \cos\text{-sim}(w'_s, w'_q)$$

since these models are optimized for this comparison function (Luong et al., 2015; Lample et al., 2018; Levy et al., 2017). Document aggregation scoring is handled identically to our SECLR models (see Section 4.2).

MT+IR. We also compare to a pipeline of NMT (Niu et al., 2018) with monolingual IR and a pipeline of SMT³ with monolingual IR. Both MT systems are trained on the same parallel sentence

³We used Moses (Koehn et al., 2007) and KenLM for the language model (Heafield, 2011).

Lang.	Analysis			Dev			Eval		
	#Q	#T	#S	#Q	#T	#S	#Q	#T	#S
Somali	300	338	142	300	482	213	1300	10717	4642
Swahili	300	316	155	300	449	217	1300	10435	4310
Tagalog	300	291	171	300	460	244	/	/	/

Table 3: MATERIAL dataset statistics: “#Q” refers to the number of queries; “#T” refers to the number of text documents; “#S” refers to the number of speech documents. There is no Tagalog Eval dataset.

Method	Somali						Swahili					
	Analysis		Dev		Eval		Analysis		Dev		Eval	
	T	S	T	S	T	S	T	S	T	S	T	S
Bivec	19.6	16.2	15.0	12.0	4.2	4.5	23.9	22.7	21.9	21.6	6.2	4.8
SID-SGNS	25.5	24.3	22.2	16.0	10.2	9.1	38.8	36.3	33.7	30.3	16.2	13.6
MUSE	9.9	9.9	10.3	16.5	1.9	2.0	27.8	24.5	27.3	28.8	9.5	8.1
NMT+IR	18.8	12.5	21.1	13.4	9.4	8.4	23.7	24.9	26.8	26.7	15.3	11.4
SMT+IR	17.4	11.2	19.1	16.8	9.1	8.3	25.5	28.6	27.1	25.2	15.4	13.3
PSQ	27.0	16.6	25.0	20.7	11.1	8.6	39.0	36.6	38.0	38.6	20.4	13.8
XLM-R	13.9	11.0	10.7	12.4	2.3	2.9	23.3	29.0	20.0	29.7	6.2	7.5
SECLR	27.8	24.4	23.0	17.4	7.7	7.4	43.8	37.9	40.3	38.1	16.0	13.1
SECLR-RT	35.4[†]	28.4	29.5	22.0	13.1[†]	11.2[†]	48.3[†]	48.1[†]	39.6	45.4	22.7[†]	17.7[†]

Table 4: Document-level MAP scores for text (T) and speech (S) for Somali and Swahili. [†] indicates significance at the $p = 0.01$ level between SECLR-RT and the best baseline.

Method	Analysis		Dev	
	T	S	T	S
Bivec	36.7	41.4	39.6	26.9
SID-SGNS	44.6	43.9	40.9	41.7
MUSE	27.4	26.5	26.0	16.5
NMT+IR	37.7	42.3	32.6	37.5
SMT+IR	44.4	52.7	39.3	35.3
PSQ	51.6	55.0	52.7	44.7
SECLR	46.7	45.0	49.3	33.9
SECLR-RT	61.1	55.5	59.0	45.7

Table 5: Document-level MAP scores for text (T) and speech (S) for Tagalog.

data as our SECLR models. The 1-best output from each MT system is then scored with Indri (Strohman et al., 2005) to obtain relevance scores. Details of NMT and SMT systems are included in Appendix C.2.

PSQ. To implement the PSQ model of Darwish and Oard (2003), we use the same alignment matrix as in rationale training (see Section 3.3) ex-

cept that here we normalize the matrix such that $\forall s \in \mathcal{V}_D, \sum_{q \in \mathcal{V}_Q} A_{q,s} = 1$. Additionally, we embed the PSQ scores into a two-state hidden Markov model which smooths the raw PSQ scores with a background unigram language model (Xu and Weischedel, 2000). The PSQ model scores each sentence and then aggregates the scores to document level as in Section 4.2.

Multilingual XLM-RoBERTa. We compare our model to the cross-lingual model XLM-RoBERTa (Conneau et al., 2020), which in previous research has been shown to have better performance on low-resource languages than multilingual BERT (Devlin et al., 2019). We use the Hugging Face implementation (Wolf et al., 2019) of XLM-RoBERTa (Base). We fine-tuned the model on the same augmented dataset of labeled query-sentence pairs as the SECLR models, but we apply the XLM-RoBERTa tokenizer before feeding examples to the model. We fine-tuned the model for four epochs using an AdamW optimizer (Loshchilov and Hutter, 2019) with learning rate 2×10^{-5} . Since XLM-RoBERTa is pretrained on Somali and Swahili but not Tagalog, we only compare our models to XLM-RoBERTa on Somali and Swahili.

5 Results and Discussion

We report Mean Average Precision (MAP) of our main experiment in Table 4 (SO & SW) and Table 5 (TL). Overall, we see that SECLR-RT consistently outperforms the other baselines in 15 out of 16 settings, and in the one case where it is not the best (SW Dev text), SECLR is the best. SECLR-RT is statistically significantly better than the best baseline on all Eval partitions.⁴ Since Analysis/Dev are relatively small, only three out of 12 Analysis/Dev settings are significant. The differences between SECLR and SECLR-RT can be quite large (e.g., as large as 70.4% relative improvement on SO Eval text), suggesting that the rationale training provides a crucial learning signal to the model.

Bivec and MUSE under-perform both of our model variants across all test conditions, suggesting that for the sentence selection task the relevance classification objective is more important than learning monolingual distributional signals. Curiously, SID-SGNS is quite competitive with SECLR, beating it on SO and SW Eval (both modalities) and TL Dev speech (five out of 16 test conditions) and is competitive with the other baselines. Again, the rationale training proves more effective as SID-SGNS never surpasses SECLR-RT.

While MT+IR is a competitive baseline, it is consistently outperformed by PSQ across all test conditions, suggesting that in low-resource settings it is not necessary to perform full translation to achieve good sentence selection performance. SMT, PSQ, and SECLR-RT all make use of the same word-alignment information but only SMT generates translations, adding additional evidence to this claim. PSQ and SECLR are close in performance on Analysis and Dev sets with SECLR eking out a slight advantage on seven of 12 Analysis/Dev set conditions.

On the larger Eval partitions, it becomes clearer that PSQ is superior to SECLR, suggesting that the relevance classification objective is not as informative as word alignment information. The relevance classification and trained rationale objectives capture slightly different information it seems; SECLR-RT, which uses both, out-performs PSQ across all 16 test conditions.

6 Training Data Ablation Study

In Section 5, we have shown that SECLR-RT consistently out-performs all baselines across all languages. Since this work targets cross-language sentence selection in a low-resource setting, we perform a training data ablation study to understand how training data size affects effectiveness.

We performed the ablation study for our two models SECLR and SECLR-RT, and the two strongest baseline methods PSQ and SID-SGNS. To simulate further the scenario of data scarcity, we sub-sampled our parallel corpus uniformly at random for 5%, 10%, 25%, 50% of the sentence pairs of the original corpus. Each sentence pair in the parallel corpus is sampled with equal probability regardless of sentence length. For consistency, for each sample size, the same sampled parallel corpus is used across all models. The word alignment probability matrix used by PSQ and SECLR-RT is generated from the same sampled corpus. Since we tune the vocabulary size on the Dev set, for fair comparison we only report MAP scores on the Analysis and Eval sets.

We plot MAP scores of the four models as a function of percentage of data sampled in Figure 1. Overall, we see that SECLR-RT consistently out-performs other baselines across all sample sizes in 9 out of 10 settings, and in the one case where it does not yield consistent improvement (Tagalog Analysis speech), SECLR-RT achieves comparable performance to PSQ.

In the low-resource setting when the sample size is 5% or 10%, SECLR consistently under-performs other models, confirming our observation that SECLR is sensitive to noise and vulnerable to learning co-occurrences of word pairs that are in fact irrelevant. When the sample size is 5% or 10%, PSQ consistently achieves better performance than SID-SGNS and SECLR (although still under-performing SECLR-RT), indicating that alignment-based methods are more robust to noise and especially useful when data is extremely scarce. The fact that SECLR-RT consistently out-performs SECLR by a wide margin for small sample sizes indicates the necessity and effectiveness of incorporating alignment-based information into SECLR to improve the robustness of the model and learn more precise alignments.

⁴We use a two-tailed paired t-test with Bonferroni correction for multiple comparisons at $p < 0.01$ for all significance tests.

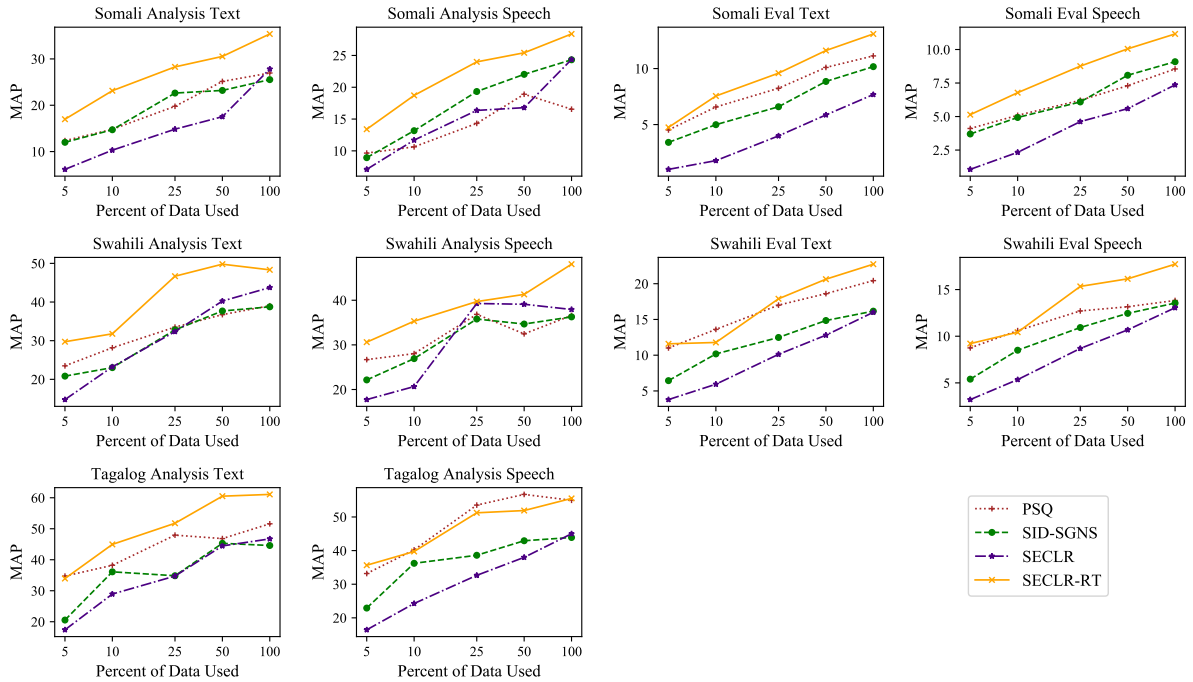


Figure 1: Ablation study results of model performances as a function of sub-sampling percentages. Note that the x-coordinate uses the log scale for better illustration of low-resource cases.

7 Alleviating the Hubness Problem

In this section, we show that by incorporating alignment information through rationale training, SECLR-RT significantly alleviates the hubness problem present in the trained cross-lingual embedding space produced by SECLR. Previous research on cross-lingual word embeddings has observed that a high-dimensional representation space with a similarity-based metric often induces a hub structure (Dinu and Baroni, 2015). Specifically, in a high-dimensional space (e.g., a cross-lingual word embedding space) defined with a pairwise similarity metric (e.g., cosine similarity), there exist a few vectors that are the nearest neighbors of many other vectors. Such vectors are referred to as “hubs.” The hub structure is problematic in IR since the hub vectors are often wrongly predicted as relevant and similar in meaning to queries that are in fact irrelevant (Radovanović et al., 2010).

Let \mathcal{V}_Q and \mathcal{V}_S be the embedding spaces for the query and sentence collection languages respectively. We define the size of the neighborhood of embeddings around $y \in \mathcal{V}_S$ as

$$N_k(y) = |\{x \in \mathcal{V}_Q | r_x(y) \leq k\}|$$

where $r_x(y)$ is the rank of y if we order \mathcal{V}_S by similarity to x from highest to lowest, and k is a

Model	Somali	Swahili	Tagalog
SECLR	29.36	54.98	43.29
SECLR-RT	6.78	14.73	11.73

Table 6: $S_{N_{10}}$ scores of SECLR and SECLR-RT respectively on Somali, Swahili and Tagalog.

positive integer. A large value of $N_k(y)$ indicates that y is similar to many $x \in \mathcal{V}_Q$, and suggests that y is a likely hub in embedding space.

Following Radovanović et al. (2010), we use $S_{N_{10}} = \mathbb{E}_{y \in \mathcal{V}_S} [(N_{10}(y) - \mu)^3 / \sigma^3]$ to measure the skewness of the distribution of N_{10} , where μ and σ refer to the mean and standard deviation of $N_{10}(y)$ respectively. Since cosine similarity is more frequently used as the similarity metric in hubness analysis, we re-train SECLR and SECLR-RT by replacing the dot product similarity metric with cosine similarity and still get performance comparable to Table 4 and Table 5.

We report $S_{N_{10}}$ scores for SECLR and SECLR-RT respectively in Table 6. We see that SECLR-RT consistently has lower $S_{N_{10}}$ value compared to SECLR on all three languages, indicating that the extra alignment information incorporated with rationale training is helpful in reducing hubness.

8 Conclusion

In this work, we presented a supervised cross-lingual embedding-based query relevance model, SECLR, for cross-language sentence selection and also applied a rationale training objective to further increase model performance. The resulting SECLR-RT model outperforms a range of baseline methods on a cross-language sentence selection task. Study of data ablation and hubness further indicate our model’s efficacy in handling low-resource settings and reducing hub structures. In future work, we hope to apply our sentence-level query relevance approach to downstream NLP tasks such as query-focused summarization and open-domain question answering.

Acknowledgements

This research is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract #FA8650-17-C-9117. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes not withstanding any copyright annotation therein.

References

- Zeynep Akkalyoncu Yilmaz, Wei Yang, Haotian Zhang, and Jimmy Lin. 2019. [Cross-Domain Modeling of Sentence-Level Evidence for Document Retrieval](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3490–3496, Hong Kong, China. Association for Computational Linguistics.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. [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)*, pages 789–798, Melbourne, Australia. Association for Computational Linguistics.
- Yujia Bao, Shiyu Chang, Mo Yu, and Regina Barzilay. 2018. [Deriving Machine Attention from Human Rationales](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1903–1913, Brussels, Belgium. Association for Computational Linguistics.
- Tal Baumel, Raphael Cohen, and Michael Elhadad. 2016. [Topic Concentration in Query Focused Summarization Datasets](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA*, pages 2573–2579. AAAI Press.
- Tal Baumel, Matan Eyal, and Michael Elhadad. 2018. [Query Focused Abstractive Summarization: Incorporating Query Relevance, Multi-Document Coverage, and Summary Length Constraints into seq2seq Models](#). *CoRR*, abs/1801.07704.
- Elizabeth Boschee, Joel Barry, Jayadev Billa, Marjorie Freedman, Thamme Gowda, Constantine Lignos, Chester Palen-Michel, Michael Pust, Banriskhem Kayang Khonglah, Srikanth Madikeri, Jonathan May, and Scott Miller. 2019. [SARAL: A Low-Resource Cross-Lingual Domain-Focused Information Retrieval System for Effective Rapid Document Triage](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 19–24, Florence, Italy. Association for Computational Linguistics.
- Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. [Reading Wikipedia to Answer Open-Domain Questions](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879, Vancouver, Canada. Association for Computational Linguistics.
- Wenhu Chen, Evgeny Matusov, Shahram Khadivi, and Jan-Thorsten Peter. 2016. [Guided Alignment Training for Topic-Aware Neural Machine Translation](#). *CoRR*, abs/1607.01628.
- Caitlin Christianson, Jason Duncan, and Boyan A. Onyshkevych. 2018. [Overview of the DARPA LORELEI Program](#). *Machine Translation*, 32(1-2):3–9.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised Cross-lingual Representation Learning at Scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Kareem Darwish and Douglas W. Oard. 2003. [Probabilistic Structured Query Methods](#). In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’03*, page 338–344, New York, NY, USA. Association for Computing Machinery.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). In *Proceedings of the 2019 Conference*

- of the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bhuwan Dhingra, Kathryn Mazaitis, and William W. Cohen. 2017. [Quasar: Datasets for Question Answering by Search and Reading](#). *CoRR*, abs/1707.03904.
- Georgiana Dinu and Marco Baroni. 2015. [Improving Zero-shot Learning by Mitigating the Hubness Problem](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Workshop Track Proceedings*.
- Yixing Fan, Jiafeng Guo, Yanyan Lan, Jun Xu, Chengxiang Zhai, and Xueqi Cheng. 2018. [Modeling Diverse Relevance Patterns in Ad-hoc Retrieval](#). In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*, pages 375–384. ACM.
- Guy Feigenblat, Haggai Roitman, Odellia Boni, and David Konopnicki. 2017. [Unsupervised Query-Focused Multi-Document Summarization Using the Cross Entropy Method](#). In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'17*, page 961–964, New York, NY, USA. Association for Computing Machinery.
- Edouard Grave, Piotr Bojanowski, Prakhar Gupta, Armand Joulin, and Tomáš Mikolov. 2018. [Learning Word Vectors for 157 Languages](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*. European Language Resources Association (ELRA).
- Surabhi Gupta, Ani Nenkova, and Dan Jurafsky. 2007. [Measuring Importance and Query Relevance in Topic-focused Multi-document Summarization](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 193–196, Prague, Czech Republic. Association for Computational Linguistics.
- Aria Haghighi, John Blitzer, John DeNero, and Dan Klein. 2009. [Better Word Alignments with Supervised ITG Models](#). In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 923–931, Suntec, Singapore. Association for Computational Linguistics.
- Kenneth Heafield. 2011. [KenLM: Faster and Smaller Language Model Queries](#). In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 187–197, Edinburgh, Scotland. Association for Computational Linguistics.
- Oana Inel, Giannis Haralabopoulos, Dan Li, Christophe Van Gysel, Zoltán Szilávik, Elena Simperl, Evangelos Kanoulas, and Lora Aroyo. 2018. [Studying Topical Relevance with Evidence-Based Crowdsourcing](#). In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management, CIKM '18*, page 1253–1262, New York, NY, USA. Association for Computing Machinery.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not Explanation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Armand Joulin, Piotr Bojanowski, Tomas Mikolov, Hervé Jégou, and Edouard Grave. 2018. [Loss in Translation: Learning Bilingual Word Mapping with a Retrieval Criterion](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2979–2984, Brussels, Belgium. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt. 2018. [Dual Conditional Cross-Entropy Filtering of Noisy Parallel Corpora](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 888–895, Belgium, Brussels. Association for Computational Linguistics.
- S. Kale, A. Kulkarni, R. Patil, Y. Haribhakta, K. Bhat-tacharjee, S. Mehta, S. Mithran, and A. Kumar. 2018. [Open-Domain Question Answering using Feature Encoded Dynamic Coattention Networks](#). In *2018 International Conference on Advances in Computing, Communications and Informatics (ICACCI)*, pages 1058–1062.
- David Kamholz, Jonathan Pool, and Susan M. Colowick. 2014. [Panlex: Building a Resource for Panlingual Lexical Translation](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014*, pages 3145–3150. European Language Resources Association (ELRA).
- Marcin Kaszkiel and Justin Zobel. 1997. [Passage Retrieval Revisited](#). In *Proceedings of the 20th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '97*, page 178–185, New York, NY, USA. Association for Computing Machinery.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A Method for Stochastic Optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran,

- Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open Source Toolkit for Statistical Machine Translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- Guillaume Lample, Alexis Conneau, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. [Word Translation without Parallel Data](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Omer Levy, Anders Søgaard, and Yoav Goldberg. 2017. [A Strong Baseline for Learning Cross-Lingual Word Embeddings from Sentence Alignments](#). *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*.
- Robert Litschko, Goran Glavaš, Ivan Vulic, and Laura Dietz. 2019. [Evaluating Resource-Lean Cross-Lingual Embedding Models in Unsupervised Retrieval](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR’19*, page 1109–1112, New York, NY, USA. Association for Computing Machinery.
- Robert Litschko, Goran Glavaš, Simone Paolo Ponzetto, and Ivan Vulić. 2018. [Unsupervised Cross-Lingual Information Retrieval Using Monolingual Data Only](#). *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*.
- Xiaoyong Liu and W. Bruce Croft. 2002. [Passage Retrieval Based on Language Models](#). In *Proceedings of the Eleventh International Conference on Information and Knowledge Management, CIKM ’02*, page 375–382, New York, NY, USA. Association for Computing Machinery.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled Weight Decay Regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Bilingual Word Representations with Monolingual Quality in Mind](#). In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 151–159, Denver, Colorado. Association for Computational Linguistics.
- Tomás Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Jian-yun Nie. 2010. [Cross-Language Information Retrieval](#). *Synthesis Lectures on Human Language Technologies*, 3:1–125.
- Xing Niu, Michael Denkowski, and Marine Carpuat. 2018. [Bi-Directional Neural Machine Translation with Synthetic Parallel Data](#). *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*.
- Franz Josef Och and Hermann Ney. 2003. [A Systematic Comparison of Various Statistical Alignment Models](#). *Computational Linguistics*, 29(1):19–51.
- Milos Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2010. [On the Existence of Obstinate Results in Vector Space Models](#). In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’10*, page 186–193, New York, NY, USA. Association for Computing Machinery.
- Anton Ragni and Mark Gales. 2018. [Automatic Speech Recognition System Development in the “Wild”](#). In *Proc. Interspeech 2018*, pages 2217–2221.
- Prajit Ramachandran, Peter Liu, and Quoc Le. 2017. [Unsupervised Pretraining for Sequence to Sequence Learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 383–391, Copenhagen, Denmark. Association for Computational Linguistics.
- Carl Rubino. 2020. [The Effect of Linguistic Parameters in CLIR Performance](#). In *Proceedings of the workshop on Cross-Language Search and Summarization of Text and Speech (CLSSTS2020)*, pages 1–6, Marseille, France. European Language Resources Association.
- Trevor Strohman, Donald Metzler, Howard Turtle, and W Bruce Croft. 2005. [Indri: A Language Model-based Search Engine for Complex Queries](#). In *Proceedings of the international conference on intelligent analysis*, volume 2, pages 2–6. Citeseer.
- Jörg Tiedemann. 2012. [Parallel Data, Tools and Interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC 2012, Istanbul, Turkey, May 23-25, 2012*, pages 2214–2218. European Language Resources Association (ELRA).
- Ivan Vulić and Marie-Francine Moens. 2015. [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, SIGIR ’15*, page 363–372, New York, NY, USA. Association for Computing Machinery.
- Courtney Wade and James Allan. 2005. [Passage Retrieval and Evaluation](#). Technical report.

- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [HuggingFace’s Transformers: State-of-the-art Natural Language Processing](#). *CoRR*, abs/1910.03771.
- Jinxi Xu and Ralph Weischedel. 2000. [Cross-Lingual Information Retrieval Using Hidden Markov Models](#). In *Proceedings of the 2000 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora: Held in Conjunction with the 38th Annual Meeting of the Association for Computational Linguistics - Volume 13, EMNLP ’00*, page 95–103, USA. Association for Computational Linguistics.
- Rabih Zbib, Lingjun Zhao, Damianos Karakos, William Hartmann, Jay DeYoung, Zhongqiang Huang, Zhuolin Jiang, Noah Rivkin, Le Zhang, Richard Schwartz, and John Makhoul. 2019. [Neural-Network Lexical Translation for Cross-Lingual IR from Text and Speech](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR’19*, page 645–654, New York, NY, USA. Association for Computing Machinery.
- Rui Zhang, Caitlin Westerfield, Sungrok Shim, Garrett Bingham, Alexander Fabbri, William Hu, Neha Verma, and Dragomir Radev. 2019. [Improving Low-Resource Cross-lingual Document Retrieval by Reranking with Deep Bilingual Representations](#). *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.

A Extra Training Dataset Details

When we train SECLR and SECLR-RT via data augmentation, we randomly split the parallel corpus into train set (96%), validation set (3%) and test set (1%). We then use the dataset augmentation technique introduced in Section 3.1 to generate positive and negative samples for each set. Augmenting the dataset upon the split corpus allows us to achieve more independence between train/validation/test set compared to splitting the dataset augmented on the entire parallel corpus. Note that we only use the validation set for early stopping but we do not tune hyperparameters with the validation set.

We preprocess the parallel corpus, the query collection and the sentence collection with the Moses toolkit (Koehn et al., 2007). The same preprocessing steps are used for all four languages (English, Somali, Swahili, Tagalog). First, we use Moses punctuation normalizer to normalize the raw text. Second, we use the Moses tokenizer to tokenize the normalized text. Finally, we remove the diacritics in the tokenized text as a cleaning step.

B Examples of Evaluation Data

In this section we demonstrate some examples from the MATERIAL dataset used for evaluation. Example queries include: “evidence”, “human rights”, “chlorine”, “academy”, “ratify”, “constitution”, “carnage” and “Kenya”. On average only 0.13% of the documents in the Eval collection are relevant to each query, which makes the task hard.

Here are two examples from Somali Analysis text. Because the documents are long, here we only include the relevant segment of a long relevant document. In the first example, the English query is “contravention” and the relevant segment of a long relevant document (translated from Somali to English by human) is “the security forces captured military equipment coming into the country illegally.” This segment is relevant to the query because of the word “illegally”.

Here is another example where the the English query is “integrity”. The relevant segment of a long relevant document (translated from Somali to English by human) is “Hargeisa (Dawan) - Ahmed Mohamed Diriye (Nana) the member of parliament who is part of the Somaliland house of representatives has accused the opposition parties (Wadani and UCID) of engaging in acts of national destruction, that undermines the existence and

sovereignty of the country of Somaliland.” This segment is relevant to the query because of the word “sovereignty”.

Since there are multiple ways to translate a word and since MT performance is relatively poor in low-resource settings, the task is far more challenging than a simple lexical match between queries and translated documents.

C Extra Experimental Details

In this section we include extra implementation and experiment details that are not included in the main paper. Information already included in the main paper are not repeated here for conciseness.

C.1 Model and Training Details

We train our SECLR and SECLR-RT models on Tesla V100 GPUs. Each model is trained on a single GPU. We report training time of SECLR and SECLR-RT on Somali, Swahili and Tagalog in Table 7.

	Somali	Swahili	Tagalog
SECLR	77	112	124
SECLR-RT	179	254	319

Table 7: Training time of SECLR and SECLR-RT on Somali, Swahili and Tagalog respectively (in minutes).

As is discussed in Section 3.2, the only trainable model parameters of SECLR and SECLR-RT are the word embedding matrices. Thus, SECLR and SECLR-RT have the same number of model parameters. We report the number of trainable parameters of both models on Somali, Swahili and Tagalog in Table 8.

	Somali	Swahili	Tagalog
# Params.	14.03M	22.31M	21.35M

Table 8: Number of trainable model parameters of SECLR/SECLR-RT on Somali, Swahili and Tagalog. “M” stands for million.

We used Mean Average Precision (MAP) as the evaluation metric in this work. We use the following implementation to compute MAP: https://trec.nist.gov/trec_eval/.

C.2 MT Baseline Details

For NMT we train bidirectional MT systems with a 6-layer Transformer architecture with model size of

Method	Somali				Swahili				Tagalog			
	Analysis		Dev		Analysis		Dev		Analysis		Dev	
	T	S	T	S	T	S	T	S	T	S	T	S
With LSTM	16.3	14.5	11.9	12.0	27.5	27.0	19.5	25.1	29.7	29.7	23.0	27.1
No LSTM	27.8	24.4	23.0	17.4	43.8	37.9	40.3	38.1	46.7	45.0	49.3	33.9

Table 9: Document-level MAP scores for text (T) and speech (S) of the SECLR model with and without LSTM.

Embed. Init.	Somali				Swahili				Tagalog			
	Analysis		Dev		Analysis		Dev		Analysis		Dev	
	T	S	T	S	T	S	T	S	T	S	T	S
Cross-lingual	35.3	27.5	31.1	23.2	48.8	41.1	42.5	41.6	56.3	51.1	53.8	45.3
Monolingual	35.4	28.4	29.5	22.0	48.3	48.1	39.6	45.4	61.1	55.5	59.0	45.7

Table 10: Document-level MAP scores for text (T) and speech (S) of the SECLR-RT model with monolingual or cross-lingual (SID-SGNS) word embedding initialization.

512, feed-forward network size of 2048, 8 attention heads, and residual connections. We adopt layer normalization and label smoothing. We tie the output weight matrix with the source and target embeddings. We use Adam optimizer with a batch size of 2048 words. We checkpoint models every 1000 updates. Training stops after 20 checkpoints without improvement. During inference, the beam size is set to 5.

Our SMT system uses the following feature functions: phrase translation model, distance-based reordering model, lexicalized reordering model, 5-gram language model on the target side, word penalty, distortion, unknown word penalty and phrase penalty.

We use backtranslation in earlier versions of MT systems. Following previous work (Niu et al., 2018), we train a bidirectional NMT model that backtranslates source or target monolingual data without an auxiliary model. This backtranslation-based model was the state-of-the-art MT model on Somali and Swahili when the above paper is published.

Later, we discover that decoder pretraining with monolingual data achieves better performance compared to backtranslation. The decoder pretraining scheme we use now is most similar to the paper by Ramachandran et al. (2017), where the authors show state-of-the-art results on the WMT English to German translation task with decoder pretraining.

There is no WMT benchmark for Somali, Swahili or Tagalog, but we use state-of-the-art techniques in our MT systems. We have also experimented with the bilingual data selection method (Junczys-Dowmunt, 2018). However, this technique does not work well, mostly because low-resource MT systems are not good enough to do scoring.

D Extra Experimental Results

In this section we include extra experimental results that are not included in the main text due to limited space.

D.1 SECLR Architecture Exploration

When we are designing the SECLR model, we experiment with adding LSTMs and using the dot product between LSTM hidden states to compute pairwise similarity between the query and the sentence. We report MAP scores of SECLR with LSTM in Table 9. Experimental results show that adding LSTMs reduces model performance consistently across all three languages. We conjecture that in low-resource settings, contextualized models create spurious correlations (Section 3.3). In fact, the XLM-RoBERTa baseline, which captures context effectively via self-attention, also underperforms our SECLR model consistently.

D.2 Word Embeddings Initialization

In our SECLR and SECLR-RT models, we initialize word embeddings with monolingual word embeddings in English, Somali, Swahili and Tagalog (Mikolov et al., 2013; Grave et al., 2018). One natural question is whether we can achieve performance improvement if we directly initialize with cross-lingual word embeddings. Because SID-SGNS out-performs both Bivec and MUSE consistently by a wide margin (Table 4 and Table 5), in this experiment we initialize SECLR-RT with the cross-lingual embeddings produced by SID-SGNS. The results of monolingual and cross-lingual embedding initialization (SID-SGNS) are shown in Table 10. We see that overall monolingual initialization slightly out-performs cross-lingual initialization. Monolingual initialization yields better performance in eight out of 12 Analysis/Dev set conditions and a MAP improvement of 1.7 points when we take the average across Analysis/Dev and all three languages.