

Instance-based Transfer Learning for Multilingual Deep Retrieval

Andrew O. Arnold
Google Research
aoa@google.com

William W. Cohen
Google Research
wcohen@google.com

Abstract

Perhaps the simplest type of multilingual transfer learning is instance-based transfer learning, in which data from the target language and the auxiliary languages are pooled, and a single model is learned from the pooled data. It is not immediately obvious when instance-based transfer learning will improve performance in this multilingual setting: for instance, a plausible conjecture is this kind of transfer learning would help only if the auxiliary languages were very similar to the target. Here we show that at large scale, this method is surprisingly effective, leading to positive transfer on all of 35 target languages we tested. We analyze this improvement and argue that the most natural explanation, namely direct vocabulary overlap between languages, only partially explains the performance gains: in fact, we demonstrate target-language improvement can occur after adding data from an auxiliary language with no vocabulary in common with the target. This surprising result is due to the effect of transitive vocabulary overlaps between pairs of auxiliary and target languages.

1 Introduction

1.1 Background and Overview

In deep retrieval, the goal is to learn to retrieve relevant documents from a large corpus of candidates, based on similarity to a query document. In multilingual transfer learning, we wish to improve performance on some task by using data from auxiliary languages to improve performance on a designated target language. Perhaps the simplest type of multilingual transfer learning is instance-based transfer learning, in which data from the target language and the auxiliary languages are pooled, and a single model is learned from the pooled data. It is not immediately obvious when instance-based transfer learning will improve performance in this multilingual setting: for instance, a plausible conjecture is

that pooling data in this way would only improve performance if the amount of auxiliary data was carefully balanced with the amount of target data, or if the auxiliary languages were carefully selected to be highly similar to the target.

In this paper we explore the behavior of instance-based transfer learning on a very large-scale task with dozens of auxiliary languages. We show that this method is surprisingly effective, leading to positive transfer on all of 35 target languages we tested, and relative improvements of up to 80%.

Analysis of this result reveals a number of influences on transfer-learning performance. Unsurprisingly, performance is improved more when the available target language data is limited, and when there is the largest overlap between target and auxiliary language vocabularies. However, analysis suggests that these effects are only a partial explanation of the effectiveness of instance-based transfer. To support this argument, we demonstrate that multilingual instance-based transfer can lead to target-language improvement after adding data from an auxiliary language with no vocabulary in common with the target: this surprising result is due to the effect of transitive vocabulary overlaps between pairs of auxiliary and target languages.

1.2 Next-Sentence Prediction as Multilingual Deep Retrieval

We examine the problem of deep retrieval on the task of next sentence prediction (NSP), where the goal is to identify the next sentence in a document from among a corpus of hundreds of millions of candidate sentences across dozens of languages, given only the current sentence as context. NSP has many important applications to areas such as question answering (Rajpurkar et al., 2018), language model training (Devlin et al., 2018), summarization (Liu et al., 2019) and conversational modeling (Vinyals and Le, 2015).

While NSP has been well studied from the perspective of sequence modeling (Ghosh et al., 2016) and binary classification (Devlin et al., 2018), we follow an alternative line of work that models the problem as an instance of deep retrieval in an extremely large output space (Logeswaran and Lee, 2018; Henderson et al., 2017; Reddi et al., 2019). Specifically, we generate a shared set of unigram and bigram features representing the *current* and *next* sentences in a NSP pair. We then train a feed-forward neural network that learns to maximize the dot-product between these consecutive sentences, as represented by the learned embedding vectors of each of the shared vocabulary’s n-gram tokens (Logeswaran and Lee, 2018).

This architecture allows us to efficiently use an extremely large vocabulary of simple unigrams and bigrams that more computationally intensive techniques do not currently permit. Using such large vocabularies, these models are able to identify and exploit tiny cross-lingual dependencies found among the many tail n-grams observed in the large unsupervised monolingual datasets available in many languages across the internet. This is in contrast to other work that tries to learn multilingual embeddings via attention (Logeswaran and Lee, 2018), shared cross-lingual embedding spaces (Chen and Cardie, 2018), cross-lingual mapping functions (Xu et al., 2018), other outside structure (Wang et al., 2013; Spohr et al., 2011) and domain knowledge (Plank and Agić, 2018; Koehn and Knight, 2002).

For language modelling, the NSP objective is interesting because it leads to difficult classification tasks for which labeled data is plentiful. The motivation for using NSP here is similar: under the deep retrieval formulation, NSP is difficult, because it requires comparison against all possible candidate sentences in the corpus, forcing the model to learn very nuanced distinctions, and potentially generating embeddings that generalize well to downstream tasks.

1.3 Instance-based Transfer Learning

As noted above, NSP is an attractive problem for unsupervised and semi-supervised transfer learning in part because of the large amount of otherwise unlabelled, yet ordered, language data available on the internet. Even for such unlabelled data, however, there is still more data available in certain languages than in others, making this problem also

an attractive setting for studying transfer learning between high and low resource languages (Pan and Yang, 2010).

We focus on *instance-based transfer learning* where we use examples drawn from a related but distinct *auxiliary* dataset to improve performance on a *target* dataset (Wang et al., 2018). However, rather than trying to align embeddings learned across languages (Søgaard et al., 2018; Nakashole, 2018; Alvarez-Melis and Jaakkola, 2018), we instead attempt the easier task of first aligning vocabulary items across languages, and then learning a single shared embedding for each vocabulary item.

Instance-based transfer vs. fine-tuning

Instance-based transfer is similar to another widely-used neural transfer method, namely *fine-tuning*. In *fine-tuning*-based transfer learning (Howard and Ruder, 2018; Huang et al., 2013; Swietojanski et al., 2012; Thomas et al., 2012), the weights of a network are trained on one (large) set of auxiliary data, and then copied into a new network where they are further adjusted using the target data. So in fine-tuning the model is optimized twice—once on the auxiliary data, and once on the target data—whereas in instance-based transfer it is optimized once, jointly, on both auxiliary and target datasets.

Fine-tuning allows additional flexibility, since one can independently decide optimization hyperparameters for each pass, but comes at some costs. In our setting, with 35 potential target languages, instance-based learning produces a single model, while fine-tuning would produce 35, one specialized for each language, and in practical settings each of these 35 different models would need to be separately stored, maintained, etc. The sequential nature of fine-tuning also means that there are many choices to make when there are multiple auxiliary languages: for instance, in learning an NSP model for Ukrainian, perhaps it is better to train first on English (the most frequent language), then Russian (a more frequent related language), and finally fine-tune on Ukrainian. In this broader setting, fully exploring the space of fine-tuned transfer models is a substantial undertaking.

In contrast, with the instance-based transfer approach we train a single joint model which shows improvements across all languages. If necessary, this efficiently computed multi-domain instance-based transfer model could be further refined by more expensive target-dependent fine-tuning; how-

ever, we leave exploration of such approaches for future work, focusing here on careful study of the more efficient instance-based method, at large scale, across many target languages.

2 Experiments

2.1 Datasets

We extract approximately 720 million next sentence pairs from a May 11, 2019 download of Wikipedia, restricting our experiments to the top 35 languages¹ accounting for 90% of the data. We split this dataset into train, development, and evaluation splits of 90%, 5% and 5% respectively. Figure 1 shows the relative size of each language in the dataset.

For each section in each article in each language we extract a pair of consecutive sentences if both sentences have at least four words. We then use a bag-of-n-grams representation for each sentence, constructing a training example as the unigram and bigram features from each $\langle current, next \rangle$ sentence pair, and aggregate all unique unigrams and bigrams into a shared *vocabulary*. For the special collection *all*, we limit the vocabulary size to the top 350 million tokens, sorted by frequency.




Figure 1: Train sentence pairs per language (log scale).

These 35 languages include some pairs that are quite similar and some pairs that are quite different. To give a rough measure of this, we looked at the similarity of the vocabularies. Figure 2 shows the *Jaccard index* among the vocabularies of the languages within the corpus, a measure of vocabulary similarity defined for a pair of language vocabularies V_a and V_b as $\frac{|V_a \cap V_b|}{|V_a \cup V_b|}$. The matrix has been sorted to emphasize clusters roughly corresponding to known language groups (e.g., the Romance languages *ro-gl-ca-pt-it-fr-es*). Other interesting

¹We use the same language code abbreviations as the Wikipedia subdomains associated with each language (e.g., *en.wikipedia.com* for English), and the special code *all* for the combined dataset comprised of data from all languages. Chinese is excluded due to the lack of a suitable tokenizer.

structure observed includes the large overlap between Serbian and Serbo-Croatian (*sr-sh*); and the cluster of Cebuano and Waray (two Austronesian languages spoken in the Philippines) with Vietnamese and Swedish (*ceb-war-vi-sv*)².




Figure 2: Jaccard index of vocabularies across languages, with blocks highlighted for Austronesian/Lsjbot (green), Serbian/Serbo-Croatian (purple), and Romance (red) languages.

2.2 Method

Using the architecture described in Section 1.2, we train two different types of models: a **per-language** model trained on a corpus containing only sentence pairs obtained from a particular language’s Wikipedia subdomain; and a **combined** model trained on a corpus containing the aggregated training examples across all languages. We train one per-language model for each language, along with a single combined model, resulting in 36 models total. We evaluate each per-language model’s performance on that language’s held-out evaluation data, and we evaluate the single combined model’s performance likewise on each individual language’s evaluation data, resulting in 35 pairs of evaluation data. We evaluate the performance of these models using a sampled **recall@k** metric defined as the percent of query sentences for which the correct next sentence is retrieved within the top-k results³. We define the **relative transfer**

²Most of the Cebuano and Waray articles were written by a computer program, *Lsjbot*, which has also written articles for Swedish Wikipedia, accounting for the large unexpected overlap in vocabularies.

³To determine the top-k results we combine the true next sentence, drawn from the evaluation data, with a sample of

improvement as:

$$\frac{\text{recall}@k_{combined} - \text{recall}@k_{per-language}}{\text{recall}@k_{per-language}}$$

3 Results

Figure 3 shows the relative transfer improvement for each language on the recall@1 metric. Perhaps surprisingly, we observe that all languages benefit from the addition of data from other languages (i.e., there is no observed negative transfer), with most languages being improved by over 25%.




Figure 3: Relative transfer improvement on each target language. Languages are sorted by number of training sentences, as in Figure 1.

3.1 Performance Improves for Many Mixture Ratios




Figure 4: Performance on Ukrainian data as a function of the *mixture ratio* between *target* data (Ukrainian) and *auxiliary* data (all excluding Ukrainian). Dashed line shows baseline performance using *native* Ukrainian mixture ratio of 2%.

These results are surprisingly strong for such a simple method. It is natural to ask if instance-based transfer can be improved, for instance by weighting the auxiliary data differently. In the experiments

100,000 false next sentence distractors, drawn from the training data. All 100,001 sentences are embedded in the same space and their similarity computed relative to the test query, with the top-k most similar sentences retrieved.

above, we built a *combined* model with relative language proportions unchanged from their original ratio in the raw Wikipedia data. Figure 4 shows the results for a particular language (Ukrainian) of adjusting this **mixture ratio** between the target language (Ukrainian) and auxiliary data (all languages except Ukrainian). We observe that performance increases as the amount of target data used increases up to a point, and then starts to diminish, with an optimum mixture ratio of 10 - 20%. In summary, gains are observed for a wide range of mixing ratios, but the chosen ratio is not optimal: there is a potential relative improvement of 17% over the *native* Ukrainian mixture ratio of 2% found in the original corpus, suggesting further improvement to the overall performance of the combined model could be achieved by optimizing all the mixture ratios across different languages.

Even though there are clearly additional improvements to be obtained by tuning instance-based transfer, a more fundamental question to ask is why the method works so well. The following sections explore this question further, by investigating potential explanations for when and why this instance-based transfer improvement occurs.

3.2 Sample Size: Low-resource Languages Improve More




Figure 5: Relative transfer improvement as a function of training data sample size, along with linear fit \pm one standard deviation. Green, purple and red points correspond to clusters in Figure 2, violet is Esperanto.

Figure 5 shows the relative transfer improvement for each language on the recall@1 metric, as a function of that language’s share of the training data. Intuitively, we observe that the languages with the smallest share of training data receive the largest improvement from this transfer. Interestingly, the language with the largest relative improvement is

Esperanto, a constructed language with vocabulary drawn from Romance and Germanic languages.

3.3 Direct Vocabulary Overlap




Figure 6: Absolute transfer improvement as a function of vocabulary overlap with **en**, along with linear fit \pm one standard deviation. Green, purple and red points correspond to clusters in Figure 2.

In addition to the observed relationship with inverse sample size, transfer also seems to improve when there is larger overlap between language vocabularies. Using English as a reference language, Figure 6 shows the absolute transfer improvement in recall@1 for each language as a function of that language’s vocabulary’s overlap with the English vocabulary $\frac{|V_{target} \cap V_{English}|}{|V_{target}|}$. The larger the overlap, the larger the benefit from transfer. This makes sense in light of the model architecture, as it is the presence of overlapping tokens from different languages in the same training sentence pairs that allows information learned on one language’s data to impact another language. Combined with the significant joint overlap among language vocabularies shown in Figure 2, this helps demonstrate why transfer can occur even between seemingly unrelated languages. Even accounting for the observed effect of sample size and direct vocabulary overlap, however, there is still a significant amount of unexplained variance in the observed transfer improvement across languages, suggesting some other effect may be at work.

3.4 Transitive Vocabulary Overlap

We design an experiment to test this hypothesis and determine whether transfer requires direct overlap between the vocabularies of two languages, or whether transfer can still happen even if the vocabularies of the two languages are disjoint. We train one model, **ca+en+ru**, on sentence pairs drawn

	ca+en+ru	+uk	% Change
Recall@1	10.6	11.3	+6.6%
Recall@10	20.1	21.5	+7.0%
Recall@20	24.3	25.6	+5.3%

Table 1: Transfer improvement of Catalan (**ca**) in the presence of Ukrainian (**uk**) data, even with no direct vocabulary overlap. English (**en**) and Russian (**ru**) serve as pivot languages.

from the three languages Catalan (**ca**), English (**en**) and Russian (**ru**), in the same proportion as they occur naturally in the training data, and evaluate this model’s performance on held out Catalan data. We then similarly sample sentence pairs from an auxiliary language, Ukrainian, whose vocabulary has been censored to contain no overlap with the Catalan vocabulary. Thus, there is no direct route by which the Ukrainian data can influence the embedding of the Catalan vocabulary (see Figure 7 for details of all vocabulary overlaps). And yet, as shown in Table 1, the addition of the Ukrainian data does improve the (**ca+en+ru+uk**) model’s performance on the Catalan data, despite there being no direct connection between the two languages.




Figure 7: Overlap among the **target**, **auxiliary** and **pivot** language vocabularies. Target \rightleftharpoons auxiliary overlap numbers (dotted arrows) are measured before censoring the vocabularies to remove any overlap.

While, by construction, there is no direct overlap between the target (**ca**) and auxiliary (**uk**) language vocabularies, there is indirect overlap via the two pivot languages, English and Russian. Thus, in a graphical sense, the influence of the auxiliary language is able to pass transitively through the chain of overlapping vocabularies of the pivot languages to ultimately influence and improve the target language’s performance. This effect seems to be enabled and amplified by the large sample sizes and diverse vocabularies of the pivot languages, even when the amount of auxiliary data is relatively small, and is most easily seen in code-mixing

around proper nouns and names (see Table 2 for an example).

... английского названия <u>Киева</u> : в результате появилось слово <u>Київ</u> . . . ими реже, чем <u>Kiev</u> . ⁴

Table 2: Example from Russian Wikipedia demonstrating code-mixing between Russian and English names for the city Kiev (underlined). Translation: *...English name of Kiev: as a result, the word Kyiv appeared...less often than Kiev.*

4 Conclusion

We have shown that cross-language instance-based transfer learning can significantly improve performance on the next sentence prediction task, when formulated as a multilingual deep retrieval problem. We have identified both sample size and vocabulary overlap as two factors that contribute to this technique’s success, and demonstrated that transfer is possible even when there is only indirect transitive vocabulary overlap. We have also shown that varying the mixture ratio between target and auxiliary data can further improve transfer performance.

These very large-scale experiments with transfer between multiple languages have uncovered some regularities not easily seen in smaller-scale experiments involving only two or three languages. For instance, the results on direct vocabulary transfer show that transfer might be further improved by increasing the vocabulary overlap among languages, using features such as byte, character, subword and phoneme n-grams (Nguyen and Chiang, 2017; Zhao et al., 2018; Wilson and Raaijmakers, 2008). More surprisingly, the results on indirect vocabulary transfer also suggest that researchers currently doing similar transfer experiments on a limited set of languages might see an immediate benefit from including a larger set of languages, even if they are seemingly unrelated, due to the transitive transfer effect. It is an open question if these kinds of transitive effects would also occur under the fine-tuning transfer approach.

References

David Alvarez-Melis and Tommi Jaakkola. 2018. Gromov-Wasserstein alignment of word embedding spaces. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*,

pages 1881–1890, Brussels, Belgium. Association for Computational Linguistics.

Xilun Chen and Claire Cardie. 2018. Unsupervised multilingual word embeddings. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 261–270, Brussels, Belgium. Association for Computational Linguistics.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Shalini Ghosh, Oriol Vinyals, Brian Strope, Scott Roy, Tom Dean, and Larry Heck. 2016. Contextual lstm (clstm) models for large scale nlp tasks. *arXiv preprint arXiv:1602.06291*.

Matthew Henderson, Rami Al-Rfou, Brian Strope, Yunhsuan Sung, László Lukács, Ruiqi Guo, Sanjiv Kumar, Balint Miklos, and Ray Kurzweil. 2017. Efficient natural language response suggestion for smart reply. *ArXiv e-prints*.

Jeremy Howard and Sebastian Ruder. 2018. Universal language model fine-tuning for text classification. *arXiv preprint arXiv:1801.06146*.

J. Huang, J. Li, D. Yu, L. Deng, and Y. Gong. 2013. Cross-language knowledge transfer using multilingual deep neural network with shared hidden layers. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 7304–7308.

Philipp Koehn and Kevin Knight. 2002. Learning a translation lexicon from monolingual corpora. In *Proceedings of the ACL-02 Workshop on Unsupervised Lexical Acquisition*, pages 9–16, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Jingyun Liu, Jackie CK Cheung, and Annie Louis. 2019. What comes next? extractive summarization by next-sentence prediction. *arXiv preprint arXiv:1901.03859*.

Lajanugen Logeswaran and Honglak Lee. 2018. An efficient framework for learning sentence representations. In *International Conference on Learning Representations*.

Ndapa Nakashole. 2018. NORMA: Neighborhood sensitive maps for multilingual word embeddings. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 512–522, Brussels, Belgium. Association for Computational Linguistics.

Toan Q. Nguyen and David Chiang. 2017. Transfer learning across low-resource, related languages for neural machine translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*,

- pages 296–301, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- S. J. Pan and Q. Yang. 2010. [A survey on transfer learning](#). *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Barbara Plank and Željko Agić. 2018. [Distant supervision from disparate sources for low-resource part-of-speech tagging](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 614–620, Brussels, Belgium. Association for Computational Linguistics.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Sashank J. Reddi, Satyen Kale, Felix Yu, Daniel Holtmann-Rice, Jiecao Chen, and Sanjiv Kumar. 2019. [Stochastic negative mining for learning with large output spaces](#). In *Proceedings of Machine Learning Research*, volume 89 of *Proceedings of Machine Learning Research*, pages 1940–1949. PMLR.
- Anders Søgaard, Sebastian Ruder, and Ivan Vulić. 2018. [On the limitations of unsupervised bilingual dictionary induction](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 778–788, Melbourne, Australia. Association for Computational Linguistics.
- Dennis Spohr, Laura Hollink, and Philipp Cimiano. 2011. A machine learning approach to multilingual and cross-lingual ontology matching. In *The Semantic Web – ISWC 2011*, pages 665–680, Berlin, Heidelberg. Springer Berlin Heidelberg.
- P. Swietojanski, A. Ghoshal, and S. Renals. 2012. [Unsupervised cross-lingual knowledge transfer in dnn-based lvcsr](#). In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pages 246–251.
- S. Thomas, S. Ganapathy, and H. Hermansky. 2012. [Multilingual mlp features for low-resource lvcsr systems](#). In *2012 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 4269–4272.
- Oriol Vinyals and Quoc Le. 2015. A neural conversational model. *arXiv preprint arXiv:1506.05869*.
- Tianyang Wang, Jun Huan, and Michelle Zhu. 2018. Instance-based deep transfer learning. *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 367–375.
- Zhigang Wang, Zhixing Li, Juanzi Li, Jie Tang, and Jeff Z. Pan. 2013. [Transfer learning based cross-lingual knowledge extraction for Wikipedia](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 641–650, Sofia, Bulgaria. Association for Computational Linguistics.
- Theresa Wilson and Stephan Raaijmakers. 2008. Comparing word, character, and phoneme n-grams for subjective utterance recognition. In *Ninth Annual Conference of the International Speech Communication Association*.
- Ruochen Xu, Yiming Yang, Naoki Otani, and Yuexin Wu. 2018. [Unsupervised cross-lingual transfer of word embedding spaces](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2465–2474, Brussels, Belgium. Association for Computational Linguistics.
- Jinman Zhao, Sidharth Mudgal, and Yingyu Liang. 2018. [Generalizing word embeddings using bag of subwords](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 601–606, Brussels, Belgium. Association for Computational Linguistics.