MODEL SELECTION FOR CROSS-LINGUAL TRANSFER USING A LEARNED SCORING FUNCTION

Anonymous authors
Paper under double-blind review

ABSTRACT

Transformers that are pre-trained on multilingual text corpora, such as, mBERT and XLM-RoBERTa, have achieved impressive cross-lingual transfer learning results. In the zero-shot cross-lingual transfer setting, only English training data is assumed, and the fine-tuned model is evaluated on another target language. No target-language validation data is assumed in this setting, however substantial variance has been observed in target language performance between different fine-tuning runs. Prior work has relied on English validation/development data to select among models that are fine-tuned with different learning rates, number of steps and other hyperparameters, often resulting in suboptimal choices. To address this challenge, we propose a meta-learning approach to model selection that uses the fine-tuned model's own internal representations to predict its cross-lingual capabilities. In extensive experiments we find that our approach consistently selects better models than English validation data across five languages and five well-studied NLP tasks, achieving results that are comparable to small amounts of target language development data.[1]

1 INTRODUCTION

Pre-trained Transformers (Vaswani et al., 2017; Devlin et al., 2019) have lead to state-of-the-art results on a wide range of NLP tasks, for example, named entity recognition, relation extraction and question answering, often approaching human inter-rater agreement (Joshi et al., 2020). These models have also been demonstrated to learn effective cross-lingual representations, even without access to parallel text or bilingual lexicons (Wu & Dredze, 2019; Pires et al., 2019). Multilingual pre-trained Transformers, such as mBERT and XLM-RoBERTa (Conneau et al., 2019), support surprisingly effective zero-shot cross-lingual transfer, where training and development data are only assumed in a high resource source language (e.g. English), and performance is evaluated on another target language.

Because no target language annotations are assumed in this setting, source language data is typically used to select among models that are fine-tuned with different hyperparameters and random seeds. However, recent work has shown that English dev accuracy does not always correlate well with target language performance (Keung et al., 2020). In this paper, we propose an alternative strategy for model selection in a zero-shot setting. Our approach, dubbed Learned Model Selection (LMS), learns a function that scores the compatibility between a fine-tuned multilingual transformer, and a target language. The compatibility score is calculated based on features of the multilingual model's learned representations and the target language. A model's features are computed based on its own internal representations; this is done by aggregating representations over an unlabeled target language text corpus. These model-specific features capture information about how the cross-lingual representations transfer to the target language after fine-tuning on source language data. In addition to model-specific representations, we also make use of learned language embeddings from the `lang2vec` package (Malaviya et al., 2017),[2] which have been shown to encode typological information, for example, whether a language has prepositions or postpositions. To measure compatibility between a multilingual model's fine-tuned representations and a target language, the model- and language- specific representations are combined in a bilinear layer. Parameters of the scoring

[1] We will make our code and data available on publication.

[2] `https://github.com/antonisa/lang2vec`

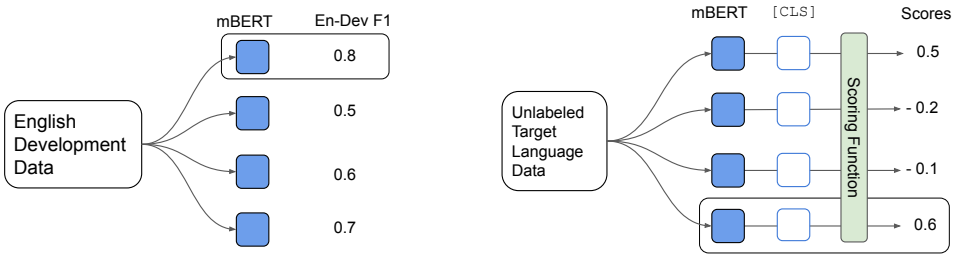

(a) English Development Selection

(b) Learned Model Selection (LMS)

Figure 1: An illustration of our approach to select the best model for zero-shot cross-lingual transfer. (a) Prior works select the best model using source language development data. (b) LMS: A learned function scores fine-tuned models based on their hidden layer representations when encoding unlabeled target language data.

function are optimized to minimize a pairwise ranking loss on a set of held-out models, where the gold ranking is calculated using standard performance metrics, such as accuracy or F_1, on a set of pivot languages (not including the target language). LMS does not rely on any annotated data in the target language for meta-learning or hyperparameter tuning, yet it is effective in learning to predict whether a multilingual model's representations are a good match for a specific target language.

In experiments on five well-studied NLP tasks (part of speech tagging, named entity recognition, question answering, relation extraction and event argument role labeling), we find LMS consistently selects models with better target-language performance than those chosen using English dev data. Appendix A.4 demonstrates that our framework supports multi-task learning, which can be helpful in settings where some target-language annotations are available, but not for the desired task. Finally, we show that LMS generalizes to both mBERT and XLM-RoBERTa in Appendix A.3.

2 Background: Zero-Shot Cross Lingual Transfer

The zero-shot setting considered in this paper works as follows. A transformer language model is first pre-trained using a standard masked language model objective. The only difference from the monolingual approach to contextual word representations (Peters et al., 2018; Devlin et al., 2019) is the pre-training corpus, which contains text written in multiple languages. For example, mBERT is trained on text written in 104 languages from Wikipedia. After pre-training on a multilingual corpus, the resulting transformer encodes language-independent representations that support surprisingly effective cross-lingual transfer, simply by fine-tuning the pre-trained parameters using English training data. For example, after fine-tuning mBERT using the English portion of the CoNLL Named Entity Recognition dataset, the resulting model can be used to perform inference directly on Spanish text, achieving an F_1 score around 75, and outperforming prior work using cross-lingual word embeddings (Xie et al., 2018; Mikolov et al., 2013). A challenge with this approach, however, is that there is a relatively high variance in this performance across training runs. Although the mean F_1 score on Spanish is 75, the performance of 60 models fine-tuned with different learning rates and random seeds ranges from around 70 F_1 to 78 (Figure 3). In zero-shot learning, no validation/development data is assumed in the target language, motivating the need for a meta-learning approach to model selection.

3 Ranking Model Compatibility with a Target Language

Given a set of multilingual BERT-based models, $M = m_1, m_2, ..., m_n$ that are fine-tuned on an English training set using different hyperparameters and random seeds, our goal is to select the model that performs the best on a target language, l_{target}. Our approach (LMS) learns to rank a set of models based on two sources of information: (1) the models' own internal representations, and (2) `lang2vec` representations of the target language (Malaviya et al., 2017). By integrating a model's

target language representations together with `lang2vec` embeddings, our approach predicts which model will perform best.

We adopt a pairwise approach to learning to rank (Burges et al., 2005; Köppel et al., 2019). The learned ranking is computed using a scoring function, $s(m, l) = f(g_{\text{mBERT}}(m), g_{\texttt{lang2vec}}(l))$, where $g_{\text{mBERT}}(m)$ is a feature vector for model m, which is computed from the model's own hidden state representations, and $g_{\texttt{lang2vec}}(l)$ is the `lang2vec` representation of language l. The model and language features are each passed through a feed-forward neural network and then combined using a bilinear layer to calculate a final score as follows:

$$\begin{aligned} s(m, l) &= f(g_{\text{mBERT}}(m), g_{\texttt{lang2vec}}(l)) \\ &= \text{FFNN}(g_{\text{mBERT}}(m))^T W_{\text{bi}} \text{FFNN}(g_{\texttt{lang2vec}}(l)) \end{aligned}$$

Using the above score, we can represent the probability that model m_i performs better than m_j on language l:

$$P(m_i \rhd_l m_j) = \sigma(s(m_i, l) - s(m_j, l))$$

where $\sigma(\cdot)$ is the sigmoid function. To tune the parameters of the scoring function, which include the feed-forward and bilinear layers, we minimize cross-entropy loss:

$$C = \sum_{l \in L \setminus \{l_{\text{target}}\}} \sum_{m_i \in M} \sum_{m_j \in M} -C_{m_i, m_j, l} \tag{1}$$

where

$$C_{m_i, m_j, l} = 1[m_i \rhd_l m_j] \log P(m_i \rhd_l m_j) + 1[m_j \rhd_l m_i] \log P(m_j \rhd_l m_i)$$

Here $1[m_j \rhd_l m_i]$ is an indicator function that has the value 1 if m_j outperforms m_i, as evaluated using labeled development data in language l.

The first sum in Equation 1 ranges over all languages where development data is available, excluding the target language. After tuning parameters to minimize cross-entropy loss on these languages, the models are ranked based on their scores for the target language, and the highest scoring model, $\hat{m} = \arg\max_m s(m, l_{\text{target}})$, is selected.

4 Tasks and Datasets

Section 3 presented an approach to selecting a model for zero-shot cross-lingual transfer. To evaluate its effectiveness, a set of fine-tuned models is split into meta-train/dev/test sets. We then use models in the meta-train set to tune parameters of the scoring function, $s(\cdot)$, develop input features with models in the meta-dev set (see §7 for details), and report our main results using models in the meta-test set in Table 2.

We experiment with five well-studied NLP tasks in the zero-shot transfer setting: named entity recognition (NER), part-of-speech (POS) tagging, question answering (QA), relation extraction (RE), and event-argument role labeling (ARL). Labeled training data for each task is assumed in English and trained models are evaluated on the target language. The task/language combinations covered by our experiments are summarized in Table 1.

The following subsections describe the tasks and datasets used in our experiments; more details are presented in Appendix A.1.

4.1 Named Entity Recognition

We use data from the CoNLL 2002 and 2003 (Sang, 2002; Sang & Meulder, 2003) shared tasks (four target languages: en, de, es, and nl), in addition to a Chinese NER dataset (Levow, 2006), following Wu & Dredze (2019). The data is tagged using the BIO tagging scheme with four types of named entities. We formulate this as a tagging task and add a linear classification layer with a softmax function to obtain word-level predictions. Because the labels are at the word level, but mBERT operates on the subword level, representation of the first subword is used to make word-level predictions, following (Devlin et al., 2019). The CoNLL evaluation script is used to compute F1 scores.

	ar	de	es	nl	zh
NER		✓	✓	✓	✓
POS	✓	✓	✓	✓	✓
QA	✓	✓		✓	✓
RE	✓				✓
ARL	✓				✓

Table 1: The five target languages and tasks used in our experiments. English is used as the source language. ar: Arabic, de: German, es: Spanish, nl: Dutch, zh: Chinese.

4.2 PART-OF-SPEECH TAGGING

A subset of Universal Dependency (UD) Treebanks (v1.4) (Nivre et al., 2018) (five target languages: ar, de, es, nl, zh) is used, following the setup of Kim et al. (2017). There are eighteen POS tags in total across all languages. We formulate this as a tagging task, following the same approach as described for named entity recognition, and report the accuracy of the predicted POS tags (ACC).

4.3 QUESTION ANSWERING

The Multilingual Question Answering (MLQA) dataset (Lewis et al., 2020) is used to evaluate model selection for cross-lingual transfer in the context of question answering. We use a subset of the MLQA dataset, consisting of four target languages: ar, de, es, zh. Following the setup of Lewis et al. (2020), models are trained using English data from SQuAD v1.1 (Rajpurkar et al., 2016), and MLQA English development set is used for the En-Dev baseline. We use the standard BERT approach for question answering Devlin et al. (2019), where a (question, passage) pair is encoded using special tokens as follows: `[BOS]` Question `[SEP]` Passage `[SEP]`. The classification layer predicts both the start and end index of the answer span from the text passage. The formal evaluation script of MLQA is adopted to compute F1 scores.

4.4 RELATION EXTRACTION

Given a pair of entity mentions, RE aims to identify a relation between them and classify its type. We use the ACE 2005 corpus (Walker et al., 2006), which contains two target languages, Arabic and Chinese, using the same preprocessed dataset as Subburathinam et al. (2019). ACE contains eighteen relation types and an additional label indicating no relation (19-way classification). We adopt the best performing model from Soares et al. (2019), [ENTITY MARKERS - ENTITY START], which adds special tokens $[\texttt{E1}_{\texttt{start}}]$, $[\texttt{E1}_{\texttt{end}}]$ and $[\texttt{E2}_{\texttt{start}}]$, $[\texttt{E2}_{\texttt{end}}]$ surrounding entity mention pairs in a sentence. The modified sentence is then fed into mBERT and representations of the starting markers $[\texttt{E1}_{\texttt{start}}]$, $[\texttt{E2}_{\texttt{start}}]$ are concatenated as inputs to a linear classification layer. We follow evaluation protocols established in prior work, using gold entities as inputs to control for differences in performance due to named entity recognition errors (Subburathinam et al., 2019). A predicted relation mention is considered correct if its type and the heads of both entities match the gold data.

4.4.1 EVENT-ARGUMENT ROLE LABELING

Event-Argument Role Labeling (ARL) aims to identify event triggers and their arguments in texts. For example, in the sentence *Facebook aquired Instagram*, an event trigger is *aquired*, whose arguments are *Facebook* and *Instagram*, with roles BUYER and SELLER. Again we draw on the ACE 2005 corpus (Walker et al., 2006) (two target languages: ar, zh), using the same preprocessed dataset as Subburathinam et al. (2019). There are thirty-five role labels in ACE. Following a similar approach as was described for relation extraction, we add special markers $[\texttt{T}_{\texttt{start}}]$, $[\texttt{T}_{\texttt{end}}]$ and $[\texttt{E}_{\texttt{start}}]$, $[\texttt{E}_{\texttt{end}}]$ around the trigger and candidate entity mentions. The modified sentence is fed into mBERT, and the marker representations are concatenated as input to a classification layer that predicts the entity's role with respect to the trigger word.

5 EXPERIMENTAL DESIGN

For a downstream task with n languages $L : \{l_1, ..., l_n\}$, our goal is to select the model that performs best on a target language, $l_{\text{target}} \in L$. We assume the available resources are English training and development data, in addition to pivot language dev data, where pivot languages are all those excluding the target language, $l_{\text{pivot}} \in L \setminus l_{\text{target}}$. First, a set of mBERT models, M, are fine-tuned on an English training set using different hyperparameters and random seeds and shuffled into meta-train/dev/test sets. We then evaluate each model, m_i, on the pivot languages' dev sets to calculate a gold ranking, $\triangleright_l$, that is used in the cross-entropy loss (Equation 1). Model-specific features are then extracted from the fine-tuned mBERTs, by feeding unlabeled pivot-language text as input.

Development and Evaluation mBERT models in the meta-dev set are used to experiment with different model and language features. Target language dev data is used to experiment with a handful of model and language features (see §7 for details).

Evaluation is performed using the meta-test set. For each target language, we rank models using the learned scoring function, select the highest scoring model, and report results in Tables 2.

5.1 BASELINES

En-Dev is our main baseline following standard practice for model selection in the zero-shot setting (Wu & Dredze, 2019; Pires et al., 2019). In addition, we compare our results with models selected using 100 sentences annotated in the target language to understand how our approach compares to the more costly alternative of annotating a small amount of target language development data. Finally, a Full-Target that picks the best model based on target language development data is adopted. These baselines and oracles are summarized below:

- En-Dev (baseline): choose an mBERT with the English dev set.
- 100-Target (oracle): choose an mBERT with 100 target language dev set instances.
- Full-Target (oracle): choose an mBERT with the full target language dev set.

5.2 HYPERPARAMETERS AND OTHER SETTINGS

To train the scoring function, $s(\cdot)$, we use Adam (Kingma & Ba, 2015), and select the batch size among $\{16, 32, 64, 128\}$, learning rate λ among $\{1\times10^{-4}, 5\times10^{-5}, 1\times10^{-5}, 5\times10^{-6}, 1\times10^{-6}\}$, and train for $\{3\}$ epochs. The scoring function, $s(\cdot)$, contains a 2-layer FFNN with 1024 hidden units and ReLU activation (Glorot et al., 2011). The base cased mBERT has 179M parameters and a vocabulary of around 120k wordpieces. Both the pre-trained transformer layers and task-specific layers are fine-tuned using Adam, with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and an L2 weight decay of 0.01. Model candidates are fine-tuned with varying learning rates and number of epochs with the following settings: learning rate $\in \{3\times10^{-5}, 5\times10^{-5}, 7\times10^{-5}\}$; number of epochs $\in \{3, 4, 5, 6\}$; batch size $\in \{32\}$. 240 mBERT models are fine-tuned and then split into meta- train/dev/test sets (120/60/60). All models are trained on a RTX 2080 Ti.

6 EVALUATION

Below we report model selection results on mBERTs in the meta-test set for each of the five tasks.

NER As illustrated in Table 2, our method selects models with a higher F1 score than En-Dev except in the case of German. In addition, it outperforms model selection using small amounts of target-language annotations (100-Target) on Dutch and selects a model that performs as well on Spanish. We use (Wu & Dredze, 2019; Pires et al., 2019) as references for zero-shot cross-lingual transfer with mBERT.

POS Table 2 presents POS accuracies on the test set, using various approaches to model selection for the five target languages. Our method selects a model with higher accuracy compared to both En-Dev and 100-Target across all five languages.

QA In Table 2, our method selects a model with higher F1 across four languages. In addition, it outperforms 100-Target on Spanish.

Methods	NER (F_1)				POS (Accuracy)					QA (F_1)				RE (F_1)		ARL (F_1)	
	de	es	nl	zh	ar	de	es	nl	zh	ar	de	es	zh	ar	zh	ar	zh
Pires et al. (2019)	69.74	73.59	77.36	-	-	-	-	-	-	-	-	-	-	-	-	-	-
Wu & Dredze (2019)	69.56	74.96	77.57	51.90	-	89.8	85.2	75.9	-	-	-	-	-	-	-	-	-
Lewis et al. (2020)	-	-	-	-	-	-	-	-	-	45.7	57.9	64.3	57.5	-	-	-	-
Subburathinam et al. (2019)	-	-	-	-	-	-	-	-	-	-	-	-	-	39.43	32.74	16.48	23.49
Model Selection																	
En-Dev	**69.89**	74.61	78.74	54.90	49.66	89.27	84.84	75.68	66.94	47.72	55.27	64.92	58.03	36.10	67.68	44.11	60.96
Ours	**69.89**	**75.74**	**78.85**	**55.05**	**51.61**	**89.77**	**85.56**	**75.87**	**68.00**	**49.26**	**55.86**	**64.95**	**58.11**	**39.54**	**70.75**	**47.08**	**62.05**
100-Target (oracle)	66.74	75.74	78.67	55.36	50.56	89.45	84.82	75.54	67.25	49.36	57.12	64.51	58.14	34.68	68.20	44.11	62.52
Full-Target (oracle)	72.07	75.73	80.26	56.88	52.70	89.98	85.11	76.04	68.84	49.36	55.83	65.08	58.38	41.92	69.13	47.15	63.81

Table 2: Model scores selected based on LMS for NER, POS, QA, RE, and ARL. En-Dev/100-Target/Full-Target: model selection based on the highest F1 of English dev set/100 target language dev set examples/target language dev set. LMS: model selection based on the highest scores for the target language: $\arg\max_m s(m, l_{\text{target}})$. Underlined scores indicate SOTA results.

ARL and RE In Tables 2, our method selects models with higher F1 score compared to En-Dev. It also outperforms 100-Target across both languages. We use $\text{GCN}_{\text{ReImp}}$ as a reference in the imbalanced dataset (see Appendix A.1 for details). Reference models were selected using the English dev set.

Model Score Distributions Figure 2 visualizes the En-Dev and LMS results on the test set in the context of the score distributions of the 60 models in the meta-test set, using kernel density estimation. English development data tends to select models that perform only slightly better than average, whereas LMS does significantly better. Similar visualizations for NER, POS, and QA are presented in Appendix A.2.

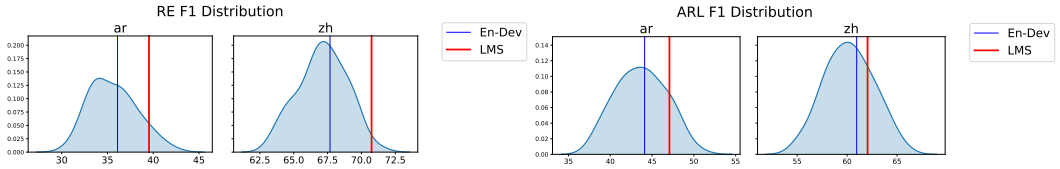

Figure 2: Model F1 score distributions for RE and ARL. Red line: LMS and blue line: En-Dev. X-axis is F1 score. Selecting models with LMS achieve better results compared to En-Dev.

7 Analysis

In Section 6, we empirically demonstrated that our learned scoring function, $s(\cdot)$, consistently selects better models than the standard approach (EN-Dev), and is comparable to small amounts of labeled target language data. Section 7 presents additional analysis of our approach, exploring the impact of various modeling choices. In addition, analysis of generalization beyond mBERT, across tasks capability, and size of training are present in Appendices A.3, A.4, and A.5.

7.1 Model and Language Features

This section explores the impact of different choices for model and language representations, during meta-learning. Four types of model features and two language embeddings are explored. We start by delineating possible choices for representations, then describe the details of our experiments, results, and the final choices used in §6.

Four model-specific features are described below. Note `[CLS]` vectors are extracted from mBERT by feeding unlabeled text as input.

- [Eng]: Averaged [CLS] vectors computed over an unlabeled English text corpus are used for both training and testing.[3]
- [Pivot]: During training, [CLS] vectors are averaged over an unlabeled text corpus in the pivot language. At test time, [CLS] embeddings are averaged over an unlabeled corpus in the target language.
- [Target]: [CLS] vectors are averaged over a text corpus in the target language (for both training and testing).
- Fusion: A linear combination of the above features. Weights on each representation are learned during training.

Two types of language embeddings are examined.

- lang2vec: 512-dimensional vectors learned by a neural network trained for typological prediction (Malaviya et al., 2017).
- syntax: 103-dimensional binary vectors, which capture syntax features from the URIEL knowledge base (Littell et al., 2017).

First, we determine the choice of model-specific features by averaging performance across both language embeddings. Table 3 reports averaged evaluation metrics for each model-specific representation across all target languages with En-Dev as a baseline.

Averaged evaluation metrics across all target languages for each language embedding are reported in Table 4. In addition to evaluating the effectiveness of each language embedding, we also experimented with a variant of our scoring function that does not include any language embeddings as input. Results are reported on mBERT models in the meta-dev set and the target languages' dev sets for all experiments in this section.

Task	*En-Dev*	[Eng]	[Pivot]	[Target]	Fusion
NER	70.45	70.64	71.18	**71.87**	70.66
POS	74.69	**75.58**	75.54	75.48	75.04
QA	56.31	56.49	**56.79**	56.68	56.63
RE	51.81	54.92	**55.57**	55.56	54.57
ARL	50.98	51.99	53.74	52.31	**54.69**
Avg	60.85	61.92	**62.60**	62.38	62.32

Table 3: Model-specific feature analysis. We use mBERT models in the meta-dev set for analysis. Each number represents average of scores across all the target languages in a particular task.

Task	*En-Dev*	lang2vec	syntax	None
NER	70.45	**71.37**	70.98	70.08
POS	74.69	**75.72**	75.36	75.20
QA	56.31	**56.81**	56.77	56.49
RE	51.81	**55.92**	55.22	52.53
ARL	50.98	53.60	**53.88**	53.14
Avg	60.85	**62.68**	62.44	61.49

Table 4: Language embedding analysis across lang2vec, syntax, and no language embedding. We use mBERT models in the meta-dev set for analysis. Each number represents average of scores across all the target languages in a particular task.

In Table 3, [PIVOT] features achieve top-2 performance in all five tasks. [Eng] and [Target] achieve mixed results, and the fusion of three features does not effectively incorporate the advantages of each representation, except in the case of ARL. Table 4 shows that lang2vec outperforms syntax for all tasks but ARL and also outperforms our approach when language embeddings are not included. Thus, lang2vec and [PIVOT] are used for all experiments in Section 6.

[3] In our experiments, sentences in the English dev set are used for this purpose (ignoring the labels).

8 RELATED WORK

Recent work has explored hyper-parameter optimization Klein et al. (2019), and model selection for a new task. `task2vec` (Achille et al., 2019) presents a meta-learning approach to selecting a pre-trained feature extractor from a library for a new visual task. More concretely, `task2vec` represents tasks in a vector space and is capable of predicting task similarities and taxonomic relations. It encodes a new task and selects the best feature extractor trained on the most similar task. Unlike `task2vec`, we select a trained model for a specific task, and we represent a trained model with model-specific features on a target language. MAML (Finn et al., 2017; Rajeswaran et al., 2019) takes another approach, pre-training a single model with a meta-loss to initialize a set of parameters that can be quickly fine-tuned for related tasks. In contrast to MAML, which is designed to support few-shot learning through better initialization of model parameters, our approach improves model selection for zero-shot cross-lingual transfer, where the training and test sets contain text in different languages.

Most relevant to our work, Xia et al. (2020) use regression methods to predict a model's performance on an NLP task. They formulate this as a regression problem based on features of the task (dataset size, average sentence length, etc.), incorporating a discrete feature to represent the choice of model. In contrast, LMS inspects a model's internal representations, thus it is suitable for predicting which out of a set of fine-tuned models will best transfer to a target language.

9 CONCLUSION

In this paper, we presented a meta-learning approach to model selection for zero-shot cross-lingual transfer. We showed that our approach improves over the standard practice of model selection using source language development data. Experiments on five well-studied NLP tasks show that by inspecting internal representations, our method consistently selects better models. LMS also achieves comparable results to the slower and more expensive alternative of annotating small amounts of target-language development data.

REFERENCES

Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless Fowlkes, Stefano Soatto, and Pietro Perona. Task2vec: Task embedding for meta-learning. In *Proceedings of the International Conference on Computer Vision*, 2019.

Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. In *Transactions of the Association for Computational Linguistics*, 2017.

Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of International Conference on Machine Learning*, 2005.

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 *Proceedings of the Association for Computational Linguistics*, 2019.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.

Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of International Conference on Machine Learning*, 2017.

Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of Artificial Intelligence and Statistics*, 2011.

Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. SpanBERT: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 2020.

M.G. Kendal. A new measure of rank correlation. *Biometrika*, (30):81–93, 1938.

Phillip Keung, Yichao Lu, Julian Salazar, and Vikas Bhardwaj. On the evaluation of contextual embeddings for zero-shot cross-lingual transfer learning. *arXiv preprint arXiv:2004.15001*, 2020. URL https://arxiv.org/pdf/2004.15001.pdf.

Joo-Kyung Kim, Young-Bum Kim, Ruhi Sarikaya, and Eric Fosler-Lussier. Cross-lingual transfer learning for POS tagging without cross-lingual resources. In *Proceedings of Emperical Methods in Natural Language Processing*, 2017.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of International Conference on Learning Representation*, 2015.

Aaron Klein, Zhenwen Dai, Frank Hutter, Neil Lawrence, and Javier Gonzalez. Meta-surrogate benchmarking for hyperparameter optimization. In *Proceedings of Advances in Neural Information Processing Systems*, 2019.

Marius Köppel, Alexander Segner, Martin Wagener, Lukas Pensel, Andreas Karwath, and Stefan Kramer. Pairwise learning to rank by neural networks revisited: Reconstruction, theoretical analysis and practical performance. *arXiv preprint arXiv:1909.02768*, 2019. URL https://arxiv.org/pdf/1909.02768.pdf.

Gina-Anne Levow. The third international chinese language processing bakeoff: Word segmentation and named entity recognition. In *Proceedings of the Fifth SIGHAN Workshop on Chinese Language Processing*, 2006.

Patrick Lewis, Barlas Oğuz, Ruty Rinott, Sebastian Riedel, and Holger Schwenk. Mlqa: Evaluating cross-lingual extractive question answering. *acl*, 2020.

Patrick Littell, David R Mortensen, Ke Lin, Katherine Kairis, Carlisle Turner, and Lori Levin. Uriel and lang2vec: Representing languages as typological, geographical, and phylogenetic vectors. In *Proceedings of the European Chapter of the Association for Computational Linguistics*, 2017.

Chaitanya Malaviya, Graham Neubig, and Patrick Littell. Learning language representations for typology prediction. In *Proceedings of Emperical Methods in Natural Language Processing*, 2017.

Tomas Mikolov, Quoc V. Le, and Ilya Sutskever. Exploiting similarities among languages for machine translation. *arXiv preprint arXiv:1309.4168*, 2013. URL https://arxiv.org/pdf/1309.4168.pdf.

Joakim Nivre, Mitchell Abrams, Željko Agić, Lars Ahrenberg, Lene Antonsen, Maria Jesus Aranzabe, Gashaw Arutie, Masayuki Asahara, Luma Ateyah, Mohammed Attia, Aitziber Atutxa, Liesbeth Augustinus, Elena Badmaeva, Miguel Ballesteros, Esha Banerjee, Sebastian Bank, Verginica Barbu Mititelu, John Bauer, Sandra Bellato, Kepa Bengoetxea, Riyaz Ahmad Bhat, Erica Biagetti, Eckhard Bick, Rogier Blokland, Victoria Bobicev, Carl Börstell, Cristina Bosco, Gosse Bouma, Sam Bowman, Adriane Boyd, Aljoscha Burchardt, Marie Candito, Bernard Caron, Gauthier Caron, Gülşen Cebiroğlu Eryiğit, Giuseppe G. A. Celano, Savas Cetin, Fabricio Chalub, Jinho Choi, Yongseok Cho, Jayeol Chun, Silvie Cinková, Aurélie Collomb, Çağrı Çöltekin, Miriam Connor, Marine Courtin, Elizabeth Davidson, Marie-Catherine de Marneffe, Valeria de Paiva, Arantza Diaz de Ilarraza, Carly Dickerson, Peter Dirix, Kaja Dobrovoljc, Timothy Dozat, Kira Droganova, Puneet Dwivedi, Marhaba Eli, Ali Elkahky, Binyam Ephrem, Tomaž Erjavec, Aline Etienne, Richárd Farkas, Hector Fernandez Alcalde, Jennifer Foster, Cláudia Freitas, Katarína Gajdošová, Daniel Galbraith, Marcos Garcia, Moa Gärdenfors, Kim Gerdes, Filip Ginter, Iakes Goenaga, Koldo Gojenola, Memduh Gökırmak, Yoav Goldberg, Xavier Gómez Guinovart, Berta Gonzáles Saavedra, Matias Grioni, Normunds Grūzītis, Bruno Guillaume, Céline Guillot-Barbance, Nizar Habash, Jan Hajič, Jan Hajič jr., Linh Hà Mỹ, Na-Rae Han, Kim Harris, Dag Haug, Barbora Hladká, Jaroslava Hlaváčová, Florinel Hociung, Petter Hohle, Jena

Hwang, Radu Ion, Elena Irimia, Tomáš Jelínek, Anders Johannsen, Fredrik Jørgensen, Hüner Kaşıkara, Sylvain Kahane, Hiroshi Kanayama, Jenna Kanerva, Tolga Kayadelen, Václava Kettnerová, Jesse Kirchner, Natalia Kotsyba, Simon Krek, Sookyoung Kwak, Veronika Laippala, Lorenzo Lambertino, Tatiana Lando, Septina Dian Larasati, Alexei Lavrentiev, John Lee, Phng Lê H`ông, Alessandro Lenci, Saran Lertpradit, Herman Leung, Cheuk Ying Li, Josie Li, Keying Li, KyungTae Lim, Nikola Ljubešić, Olga Loginova, Olga Lyashevskaya, Teresa Lynn, Vivien Macketanz, Aibek Makazhanov, Michael Mandl, Christopher Manning, Ruli Manurung, Cătălina Mărănduc, David Mareček, Katrin Marheinecke, Héctor Martínez Alonso, André Martins, Jan Mašek, Yuji Matsumoto, Ryan McDonald, Gustavo Mendonça, Niko Miekka, Anna Missilä, Cătălin Mititelu, Yusuke Miyao, Simonetta Montemagni, Amir More, Laura Moreno Romero, Shinsuke Mori, Bjartur Mortensen, Bohdan Moskalevskyi, Kadri Muischnek, Yugo Murawaki, Kaili Müürisep, Pinkey Nainwani, Juan Ignacio Navarro Horñiacek, Anna Nedoluzhko, Gunta Nešpore-Bērzkalne, Lng Nguy˜ên Thị, Huy`ên Nguy˜ên Thị Minh, Vitaly Nikolaev, Rattima Nitisaroj, Hanna Nurmi, Stina Ojala, Adédayọ̀ Olúòkun, Mai Omura, Petya Osenova, Robert Östling, Lilja Øvrelid, Niko Partanen, Elena Pascual, Marco Passarotti, Agnieszka Patejuk, Siyao Peng, Cenel-Augusto Perez, Guy Perrier, Slav Petrov, Jussi Piitulainen, Emily Pitler, Barbara Plank, Thierry Poibeau, Martin Popel, Lauma Pretkalniņa, Sophie Prévost, Prokopis Prokopidis, Adam Przepiórkowski, Tiina Puolakainen, Sampo Pyysalo, Andriela Rääbis, Alexandre Rademaker, Loganathan Ramasamy, Taraka Rama, Carlos Ramisch, Vinit Ravishankar, Livy Real, Siva Reddy, Georg Rehm, Michael Rießler, Larissa Rinaldi, Laura Rituma, Luisa Rocha, Mykhailo Romanenko, Rudolf Rosa, Davide Rovati, Valentin Roșca, Olga Rudina, Shoval Sadde, Shadi Saleh, Tanja Samardžić, Stephanie Samson, Manuela Sanguinetti, Baiba Saulīte, Yanin Sawanakunanon, Nathan Schneider, Sebastian Schuster, Djamé Seddah, Wolfgang Seeker, Mojgan Seraji, Mo Shen, Atsuko Shimada, Muh Shohibussirri, Dmitry Sichinava, Natalia Silveira, Maria Simi, Radu Simionescu, Katalin Simkó, Mária Šimková, Kiril Simov, Aaron Smith, Isabela Soares-Bastos, Antonio Stella, Milan Straka, Jana Strnadová, Alane Suhr, Umut Sulubacak, Zsolt Szántó, Dima Taji, Yuta Takahashi, Takaaki Tanaka, Isabelle Tellier, Trond Trosterud, Anna Trukhina, Reut Tsarfaty, Francis Tyers, Sumire Uematsu, Zdeňka Urešová, Larraitz Uria, Hans Uszkoreit, Sowmya Vajjala, Daniel van Niekerk, Gertjan van Noord, Viktor Varga, Veronika Vincze, Lars Wallin, Jonathan North Washington, Seyi Williams, Mats Wirén, Tsegay Woldemariam, Tak-sum Wong, Chunxiao Yan, Marat M. Yavrumyan, Zhuoran Yu, Zdeněk Žabokrtský, Amir Zeldes, Daniel Zeman, Manying Zhang, and Hanzhi Zhu. Universal dependencies 2.2, 2018. URL `http://hdl.handle.net/11234/1-2837`.

Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2018.

Telmo Pires, Eva Schlinger, and Dan Garrette. How multilingual is multilingual BERT? In *Proceedings of the Association for Computational Linguistics*, 2019.

Aravind Rajeswaran, Chelsea Finn, Sham M Kakade, and Sergey Levine. Meta-learning with implicit gradients. In *Proceedings of Advances in Neural Information Processing Systems*, 2019.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of Emperical Methods in Natural Language Processing*, 2016.

Erik F. Tjong Kim Sang. Introduction to the CoNLL-2002 shared task: Language-independent named entity recognition. In *Proceedings of International Conference on Computational Linguistics*, 2002.

Erik F. Tjong Kim Sang and Fien De Meulder. Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition. In *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2003.

Livio Baldini Soares, Nicholas FitzGerald, Jeffrey Ling, and Tom Kwiatkowski. Matching the blanks: Distributional similarity for relation learning. In *Proceedings of the Association for Computational Linguistics*, 2019.

Ananya Subburathinam, Di Lu, Heng Ji, Jonathan May, Shih-Fu Chang, Avirup Sil, and Clare Voss. Cross-lingual structure transfer for relation and event extraction. In *Proceedings of Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, 2019.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of Advances in Neural Information Processing Systems*, 2017.

Christopher Walker, Stephanie Strassel, Julie Medero, and Kazuaki Maeda. ACE 2005 multilingual training corpus. In *Linguistic Data Consortium*, 2006. LDC2006T06.

Shijie Wu and Mark Dredze. Beto, Bentz, Becas: The surprising cross-lingual effectiveness of BERT. In *Proceedings of Empirical Methods in Natural Language Processing and the International Joint Conference on Natural Language Processing*, 2019.

Mengzhou Xia, Antonios Anastasopoulos, Ruochen Xu, Yiming Yang, and Graham Neubig. Predicting performance for natural language processing tasks. In *Proceedings of the Association for Computational Linguistics*, 2020.

Jiateng Xie, Zhilin Yang, Graham Neubig, Noah A Smith, and Jaime G Carbonell. Neural cross-lingual named entity recognition with minimal resources. In *Proceedings of Emperical Methods in Natural Language Processing*, 2018.

Wei Ye, Bo Li, Rui Xie, Zhonghao Sheng, Long Chen, and Shikun Zhang. Exploiting entity BIO tag embeddings and multi-task learning for relation extraction with imbalanced data. In *Proceedings of the Association for Computational Linguistics*, 2019.

A APPENDIX

A.1 IMBALANCED DATA FOR RE AND ARL

In this section, we describe details of the imbalanced dataset for RE and ARL. Table 5 reports the statistics of the dataset and Table 6 summarizes references and baseline results.

The number of non-relation entity pairs or trigger-entity pairs (negative instances) far exceeds positive instances, which negatively affects a model's performance. However, the preprocessed dataset from (Subburathinam et al., 2019) is artificially balanced, which does not reflect a practical setting. Therefore, we create an additional imbalanced dataset using the ACE2005 corpus Walker et al. (2006), which more closely replicates the setting a model will be faced with in a real-world information extraction scenario. First, we shuffle documents into 80%/10%/10% splits for train/dev/test, then extract candidate entity-pairs from each document. For RE, the first approach in (Ye et al., 2019) is adopted to extract negative instances. Negative instances whose entity-type combination has never appeared as a positive example in the training data are filtered out. This results in a positive/negative ratio of 1:8.9 within 80,256 mention pairs for English. We repeat the same process for Chinese/Arabic and end up with a positive/negative ratio of 1:9.4/1:8.9 and 73,082/42,104 instances. For ARL, we create negative instances by pairing each trigger with every entity in a sentence. This leads to a positive/negative ratio of 2.6/2.7/2.9 and 27,823/19,338/14,376 instances in English/Chinese/Arabic. Details on the two datasets are summarized in Table 5.

As a baseline for the imbalanced dataset, we reimplement the Graph Convolutional Network (GCN) model of (Subburathinam et al., 2019) using multilingual embeddings learned by fastText (Bojanowski et al., 2017) on Wikipedia ($\text{GCN}_{\text{ReImp}}$). Tables 6 display F1 for zero-shot cross-lingual transfer in both balanced and imbalanced datasets.

Task	Lang	Train	Dev	Test	Pos/Neg
RE	en	63177	10218	6861	1:8.9
	zh	57824	7096	8162	1:9.4
	ar	32984	4482	4638	1:8.9
ARL	en	21875	3345	2603	1:2.6
	zh	15095	2226	2017	1:2.7
	ar	11587	1221	1568	1:2.9

Table 5: Statistics of the imbalanced dataset. Number of instances and the total positive/negative ratio.

	RE ($\mathbf{F}_1$)		**ARL ($\mathbf{F}_1$)**	
	ar	zh	ar	zh
Balanced Dataset				
GCN (Subburathinam et al., 2019)	58.70	42.50	61.80	59.00
$\text{GCN}_{\text{ReImp}}$	56.10	41.70	62.08	55.10
mBERT	69.08	79.56	59.27	68.00
Imbalanced Dataset				
$\text{GCN}_{\text{ReImp}}$	39.43	32.74	16.48	23.49
Model Selection (Imbalanced Dataset)				
En-Dev	36.10	67.68	44.11	60.96
LMS	**39.54**	**70.75**	**47.08**	**62.05**
100-Target	34.68	68.20	44.11	62.52
Full-Target	41.92	69.13	47.15	63.81

Table 6: F1 scores for relation extraction and argument role labeling on the test set. En-Dev/100-Target/Full-Target: model selection based on the highest F1 of English dev set/100 target language dev set examples/target language dev set. Ours: model selection based on the highest scores for the target language: $\arg\max_m s(m, l_{\text{target}})$.

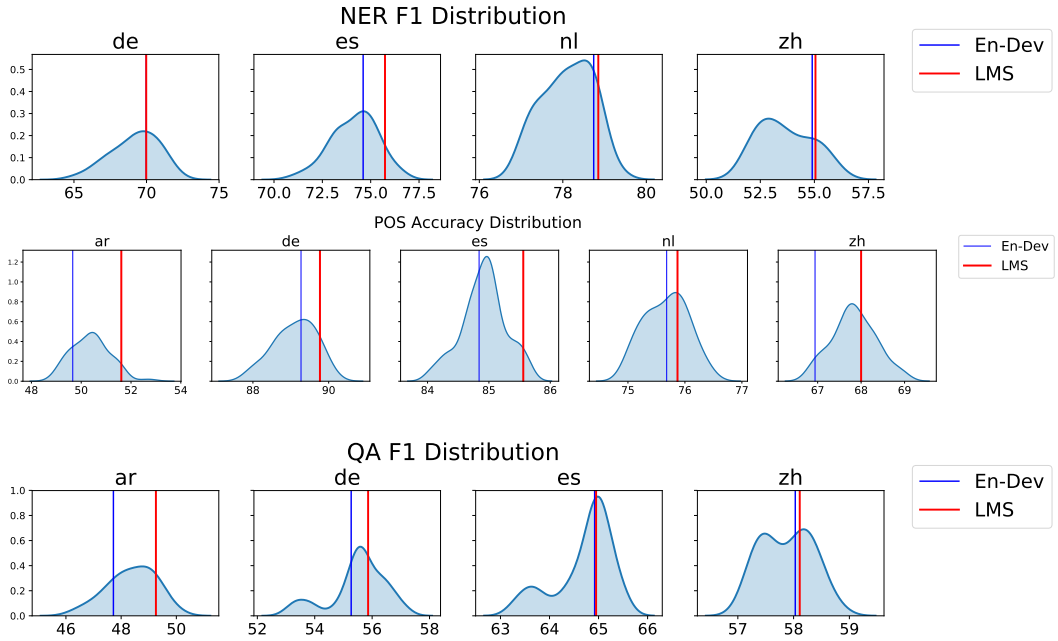

Figure 3: Model score distribution of models in the meta-test set for NER, POS, and QA. LMS (Red) and En-Dev (Blue). X-axis is F1 score for NER and QA and accuracy for POS.

A.2 Model Score Distribution

Details of model score distribution are explained and analyzed in this section. Distribution figures of NER, POS, and QA are shown in Figure 3.

To better interpret model selection results, we visualize the distribution of model scores in the meta-test set with the kernel density estimate (KDE) plot. KDE represents the data using a continuous probability density curve in one or more dimensions. LMS (Red) and En-Dev (Blue) selection results are marked in the model score distribution. Majority model scores are concentrated at the center of the distribution but the variations are large. For example, Arabic POS accuracy in Figure 3 ranging from 48 to 53, where the majority models achieve a score around 50. En-Dev fails to select a model from the middle of the distribution, which leads to a score below average models. However, LMS is able to not only select a model better than the En-Dev but also select a model that is better than the majority of candidates. In fact, En-Dev is not a good model selection strategy when it comes to linguistically different languages such as Arabic and Chinese. For Arabic POS, Arabic QA, and Chinese POS, En-Dev selects a model that performs worse than an average model from the candidates (left side of the distribution). This might suggest the model is over-fitting to English data but fails to transfer to a distanced language. Our observations motivate us to develop a better model selection method - LMS.

A.3 Does this Approach Generalize to XLM-RoBERTa?

In Section 6, we showed that our approach consistently selects better fine-tuned models than those chosen using English development data. To test the robustness of our approach with a different multilingual pre-trained transformer, we evaluate it on XLM-RoBERTa-base (Conneau et al., 2019), the same setting used for mBERT in Section 6 when testing RE and ARL with imbalance dataset.

RE In the left section of Table 7, our approach selects a model with a higher F1 score compared to En-Dev in Chinese and on par with En-Dev in Arabic.

ARL In the right section of Table 7, our approach selects a model with a higher F1 score compared to En-Dev in Arabic but performs worse on Chinese. Overall, our approach appears to be effective when used with XLM-RoBERTa.

	RE (F_1)		ARL (F_1)	
	ar	zh	ar	zh
GCN_{ReImp}	39.43	32.74	16.48	23.49
mBERT	36.10	67.68	44.11	60.96
Model Selection				
En-Dev	**40.79**	64.48	50.65	**62.73**
LMS	**40.79**	**65.11**	**52.96**	61.92
100-Target	42.33	65.38	52.90	62.12
Full-Target	44.66	65.75	53.09	62.27

Table 7: XLM-RoBERTa experiment: F1 of relation extraction and argument role labeling on the imbalanced dataset. Model selection results are based on XLM-RoBERTa-base models in the meta-test set.

A.4 CAN MULTI-TASK LEARNING HELP?

Our setting does not assume access to the labeled data in the target language for a particular task. However, labeled data in the target language may be available for a relevant auxiliary task, which could help the scoring function learn to better estimate whether a model is a good match for the target language.

To test whether an auxiliary task in the target language might help to select a better model for the target task, we fine-tune sets of mBERT models for ARL and RE. Gold rankings on the models are then computed for each language using the pivot languages' dev sets. In addition, another "silver" ranking is computed for each language using the auxiliary task. The scoring function is then trained to rank mBERT models for both tasks. To differentiate the two tasks, a variant of scoring function, $s(m, l, t)$, which concatenates a randomly initialized task embedding with the language embedding is adopted. Following Section 7, we use model selection results from mBERT models in the meta-dev set and report average target language dev-set results across Chinese and Arabic.

In Table 8, our approach is able to select a model with higher F1 score for RE. However, multi-task does not benefit ARL but still outperforms En-Dev.

Task	*En-Dev*	`([Pivot], lang2vec)`	+ Multi-task
RE	51.81	55.92	**57.31**
ARL	50.98	**53.60**	51.99

Table 8: Multi-task analysis using additional training data in the target language from another task. We use mBERT models in the meta-dev set for analysis. Model selection is based on the highest scores for the target language and target task: $\arg\max_m s(m, l_{\text{target}}, t_{\text{target}})$

A.5 HOW MANY MODELS DO WE NEED TO TRAIN LMS?

We analyze how practical LMS is in this section. In particular, we aim to answer the question of how many models do we need in order to train LMS. In fact, LMS can be trained with only six models.

LMS outperforms the En-Dev in three out of four target languages by training on only six models (Table 9). Varying the number of models in the meta-training set from 6 to 120 for QA, we see LMS perform consistently across different numbers of training sizes. It outperforms the En-Dev when trained with 72 models.

To better understand how LMS performs compared to En-Dev, we use Kendall Rank Correlation Coefficient (KRCC) (Kendal, 1938) to measure the similarity of the orderings. KRCC (-1 to 1) between two model rankings will be high when have a similar rank. In Figure 4, KRCC is calculated between En-Dev/LMS and golden model rankings for 60 models in the QA meta-test set. LMS outperforms En-Dev with only six models and the KRCC improves as more models are used to train LMS. It is interesting to see LMS has the highest KRCC in Arabic. We hypothesis this is because of the high variation of F1 score in Arabic - the standard deviation of the Arabic F1 score is 3.3 as opposed to 2.4 and 1.3 in German and Spanish.

	ar	de	es	zh
En-Dev	47.72	55.27	64.92	58.03
# *of models to learn* LMS				
6	48.95	55.21	65.28	58.27
8	48.95	55.21	65.28	58.27
72	48.95	56.03	64.95	58.27
120	49.26	55.86	64.95	58.11

Table 9: Model selection results for QA (F1). A varying number of models to train LMS from 6 to 120.

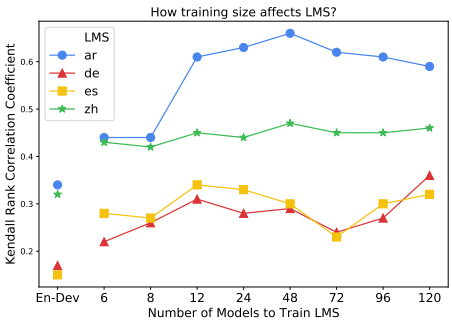

Figure 4: The similarity between En-Dev/LMS ranking and the golden ranking for models in the QA meta-test set.