

# DaLC: Domain Adaptation Learning Curve Prediction for Neural Machine Translation

Cheonbok Park<sup>†</sup>, Hantae Kim<sup>†</sup>, Ioan Calapodescu<sup>‡</sup>

Hyunchang Cho<sup>†</sup>, and Vassilina Nikoulina<sup>‡,\*</sup>

<sup>†</sup>Papago, NAVER Corp., <sup>‡</sup>NAVER LABS Europe

{cbok.park, hantae.kim, hyunchang.cho}@navercorp.com  
{ioan.calapodescu, vassilina.nikoulina}@naverlabs.com

## Abstract

Domain Adaptation (DA) of Neural Machine Translation (NMT) model often relies on a pre-trained general NMT model which is adapted to the new domain on a sample of in-domain parallel data. Without parallel data, there is no way to estimate the potential benefit of DA, nor the amount of parallel samples it would require. It is however a desirable functionality that could help MT practitioners to make an informed decision before investing resources in dataset creation. We propose a Domain adaptation Learning Curve prediction (DaLC) model that predicts prospective DA performance based on in-domain *monolingual* samples in the source language. Our model relies on the NMT encoder representations combined with various instance and corpus-level features. We demonstrate that instance-level is better able to distinguish between different domains compared to corpus-level frameworks proposed in previous studies (Xia et al., 2020; Kolachina et al., 2012). Finally, we perform in-depth analyses of the results highlighting the limitations of our approach, and provide directions for future research.

## 1 Introduction

The classical Domain Adaptation scenario (Freitag and Al-Onaizan, 2016; Luong and Manning, 2015) usually relies on an existing NMT model trained on large datasets originated from various sources. This *general* model is adapted to the new domain, with a small sample of *in-domain parallel* data, through finetuning or other DA methods. Without any parallel in-domain data, we cannot estimate the quality of *general* NMT model on the domain of interest, nor can we anticipate what will be the benefits of the DA and how much parallel in-domain data is required.

In this work, we address the problem that translation service providers may regularly face when

receiving a request for a new domain translation. In such case, a new domain is often defined by its source language samples, and the translation provider needs to invest into in-domain parallel dataset creation in order to be able to perform evaluation and Domain Adaptation of its *general* model. Current state of the art research in NMT Domain Adaptation rarely provides any insights on the amount of data required to perform Domain Adaptation depending on input domain characteristics. This is however a desirable feature that would allow to (i) estimate the data creation cost (time and money-wise) for the client requesting an adaptation to completely new domain; (ii) make an informed decision on how to distribute fixed data creation budget when there is a need to handle multiple DA simultaneously (as we demonstrate further sharing this budget equally across all domains may not be optimal). The goal of this work is to gain better insights on Domain Adaptation dynamics, and provide practical guidelines for such a real-life scenario.

Several studies address the problem of learning curve estimation models (Xia et al., 2020; Ye et al., 2021; Kolachina et al., 2012) of MT or NLP models without actually training those. This is done by training a prediction model which takes corpus-level representation  $X$  as an input and predicts the score  $y$  for this corpus as an output. In the case of MT learning curves prediction,  $X$  would correspond to the parallel data sample (used to train MT model), and  $y$  is the BLEU score achieved by an MT model trained on  $X$ . This means that each training point creation requires training a new MT model, which may become very costly if we want to create a training set of reasonable size. In this work, we propose a novel framework to perform the learning of the prediction model *at instance-level*. It can significantly decrease the cost of training samples creation, by leveraging instance-level representations.

\* Corresponding author

There are multiple factors that would impact DA learning curve: the complexity of the new domain, characteristics of the in-domain samples, baseline performance on the new domain, DA algorithm and its hyperparameters. In this work, we concentrate on the factors related to domain complexity and in-domain sample characteristics. We are particularly interested to understand how much can we get from source only in-domain sample, which corresponds to the real-life scenario.

Aharoni and Goldberg (2020) points out that the notion of domain can be fuzzy. They suggest that pretrained language models (LMs) representation could contain rich information about different domains, and have demonstrated that there might be significant overlap between different domains at instance-level. In the context of NMT Domain Adaptation, NMT encoder representations are better suited to characterize different domains, and to evaluate the difficulty of those domains for the general NMT model (Del et al., 2021).

The main contributions of this work are:

- (i) We formulate the problem of DA learning curves prediction as instance-level framework and demonstrate that instance-level representation favours fine-grained knowledge transfer across different domains thus significantly decreasing the cost of training samples creation;
- (ii) We propose a prediction model that relies on NMT encoder representations combined with a number of other instance and corpus-level characteristics, computed from the monolingual (source) in-domain sample only;
- (iii) We analyse how far we can go in DA performance prediction based on the source side information only, and outline some limitations of this constraint in Section 7.

## 2 Related work

The problem of Learning Curve prediction can be related to a number of different existing problems in Natural Language Processing (NLP) and Machine Learning. In this section, we briefly review few works that have tried to predict Learning Curves (also known as *Scaling Laws*) or domain shift for different NLP tasks. We also overview some works on Active Learning for NMT since those features can also be relevant to our task.

### Learning Curve and domain shift prediction.

There is a number of works that have attempted to predict model’s performance without actually executing (and even training) the model in different contexts. Elshahar and Gallé (2019) predicts the classifier’s performance drop under the domain shift problem; Elloumi et al. (2018) estimates the ASR performance for unseen broadcasts. Xia et al. (2020); Ye et al. (2021) study the problem of predicting a new NLP task performance based on the collection of previous observed tasks of different nature. Their proposed models are evaluated on MT among other tasks.

Closest to our work, Kolachina et al. (2012) predicts the learning curves for the Statistical Machine Translation(SMT) task. They formulate the task as a parametric function fitting problem, and infer the learning curve relying on set of features based on in-domain source and/or target sample (assuming no in-domain parallel sample is available). Most of the above-mentioned works rely on corpus-level score predictions and therefore require large amount of trained MT model instances to generate sufficient amount of training points for the predictor model.

There is very recent interest around modeling the *Scaling Laws* for a language model (Kaplan et al., 2020) or a NMT model (Gordon et al., 2021; Ghorbani et al., 2021). Similar to (Kolachina et al., 2012) these works try to derive a parametric function that would allow to make a connection between the model’s training characteristics (amount of data, parameters or compute) and model’s final performance. These models operate at corpus level and do not address the problem of Domain Adaptation (and the fact that different domains may follow different scaling laws functions).

To best of our knowledge, our work is the first attempt to specifically solve learning curves prediction in the context of NMT Domain Adaptation. We propose an instance-level framework relying on NMT encoder representations (which none of the previous work did) in combination with other features. Owing to this framework, our proposed framework requires a small amount of trained MT models with less than 10 models.

**Active learning for NMT.** Active Learning (AL) algorithms are built to select the most useful samples for improving the performance of a given model. Therefore, the criteria used by AL algorithms in NLP or MT tasks (Zhang et al., 2017; Zhao et al., 2020; Peris and Casacuberta, 2018;

Dou et al., 2020; Wang et al., 2020; Zhan et al., 2021) could also serve as discriminative features when it comes to predicting the future NMT performance. We reuse some of the scoring functions introduced by previous work as specified in Section 4.2. On the other hand, we believe that a successful Learning Curves prediction framework can help identify important features and/or samples for an AL framework.

### 3 Analysis of real Domain Adaptation Learning curves

In this section, we analyse several real learning curves for NMT Domain Adaptation to show how learning curves do behave across domains and to motivate our work. In order to perform this analysis, we first present the NMT baseline model (i.e., *general* model) we utilize, as well as datasets used for Domain Adaptation. We also discuss the NMT evaluation metrics that we can rely on when training the learning curve predictor model<sup>1</sup>.

**NMT model.** We consider two different NMT systems: English-German and German-English systems trained on WMT20 dataset (Barrault et al., 2020). We provide technical details about architecture and datasets used to train those NMT systems in the Appendix A.

**Domain Adaptation data.** We rely on the dataset released by Aharoni and Goldberg (2020). The dataset consists of train/dev/test with duplicated sentences. This dataset splits for 5 domains (Koran, IT, Medical, Law, Subtitles) from OPUS (Tiedemann, 2012) for German-English. For each domain, we create random samples  $S_{n,d}$  of size  $n$  (with  $n = 1K, 10K, 20K, 100K$ ) at  $d$  domain. Those samples are then used to train instances of Domain Adapted models  $M_{S_{n,d}}$  resulting in total 19 models<sup>2</sup>. We will further refer to the size of these different samples as *anchor points* (of the learning curve).

**Domain Adaptation.** For each *general* NMT model (*baseline*), we create a set of Domain Adapted models trained on different samples of in-domain data described previously (anchor points).

<sup>1</sup>In what follows we may refer to the *learning curve predictor model* as *predictor* to avoid the confusion with NMT model.

<sup>2</sup>We didn't train 100K-sampled DA model for *Koran* domain since it only has 20K parallel sentences in the training split

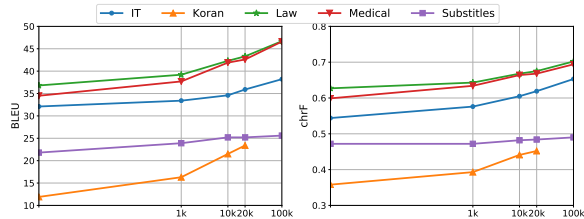


Figure 1: Domain Adaptation Learning Curves for English-German NMT model. Y-axis represents the performance of Domain adapted model (e.g., chrF or BLEU), X-axis: amount of in-domain samples used by Domain Adaptation (at log scale).

Domain Adaptation is done via finetuning only with an in-domain dataset (Appendix A provides details).

**NMT evaluation.** Our main goal is to obtain a corpus-level score that would allow us to assess the DA performance which is traditionally measured by BLEU score (Papineni et al., 2002). While BLEU score may exhibit reasonable correlation with human judgements at corpus-level, it is known to have poor correlation at instance-level. Recall that we are interested to exploit instance-level representations to favour knowledge transfer across domains, therefore we require a reliable instance-level metric to create gold annotations that the predictor model could learn from. We rely on *chrF* (Popović, 2015) score that, according to WMT 20 MT evaluation track (Mathur et al., 2020), provides reasonable correlation with human judgement for instance-level evaluations. For learning curves prediction we rely on *mean chrF* (average instance-level chrF across the whole test set) as a proxy for corpus-level score thus making connection with instance-level scores used for training of the predictor.

**Analysis of Learning Curves.** Figure 1 reports how BLEU and *mean chrF* scores progress with amount of in-domain samples used for Domain Adaptation (English-German). First, we note that *mean chrF* metric exhibits the same behaviour as BLEU when tracing the Learning Curve<sup>3</sup>.

Second, we note the difference in a Domain Adaptation progress for each of domains. For instance, *Koran* learning curve is quite steep, while *Subtitles* domain improves very slightly and reaches a plateau performance at already 1K Do-

<sup>3</sup>Table 5 in Appendix A provides BLEU/chrF for German-English and confirms the above observation. All the evaluations are performed with SacreBLEU toolkit (Post, 2018).

main Adaptation anchor point. We also note that this behaviour isn't necessarily linked to the performance of the *baseline* model on those domains: *Law* and *Medical* domains achieve quite high scores when translated with the *baseline* NMT model (tables 5, 6 in the appendix), but they keep improving as in-domain data grows. *Subtitles* on the other hand has lower *baseline* BLEU and *mean chrF* scores and Domain Adaptation with growing in-domain data doesn't seem to help much. Such behaviour probably reveals the limitations of the dataset, and/or a Domain Adaptation method used for this domain that would merit further investigations. The problem that we want to address in our work is whether it is possible to forecast such behaviour in advance, and how far we can go in this task with source only in-domain data sample.

## 4 Approach

In this section, we first formalize the DA learning curve predictor. Then, we describe the representations and our model for the problem.

### 4.1 Problem setup

We are given a *baseline model*  $M_G$  (trained on general corpus  $G$ ), input sentence  $x$  and a new domain  $d$  defined by its sample  $S_d$ . The *DA learning curve predictor* can be modeled as a scoring function  $g_\theta$  which depends on the instance-level representation  $\phi(x)$ , and corpus-level representation  $\xi(S_d^s)$ <sup>4</sup>.

Model  $M_{S_d}$  is an NMT model obtained by adapting  $M_G$  to the new domain (i.e.,  $d$  on the in-domain sample of parallel sentences  $S_d$ ). Learning of *DA learning curve predictor* can be done by regressing the actual scoring function  $y = s(x, M_{S_d})$  that provides translation quality score for an input sentence (from the test domain)  $x$  translated with  $M_{S_d}$ .  $s(x, M_{S_d})$  refers to *chrF* score as discussed in Section 3.

The learning objective is then formulated as

$$\min_{\theta} \sum_{d \in D} \sum_{x \in T_d} (s(x, M_{S_d}) - g_{\theta}(\phi(x), \xi(S_d^s)))^2$$

where  $D$  and  $T_d$  are a set of training domains<sup>5</sup> and training sentences of  $d$  domain<sup>6</sup>, respectively.

<sup>4</sup> $S_d^s$  denotes source side of in-domain sample  $S_d$  since we restrict ourselves to the case where we can only access source side of in-domain samples.

<sup>5</sup>We consider the scenario where the test domain is not known during the predictor learning phrase and disjoint form

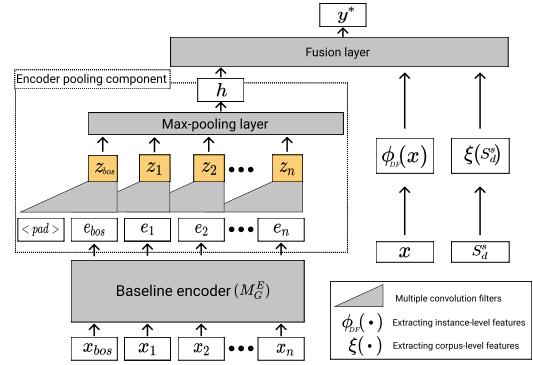


Figure 2: Overview of the learning curve predictor model, where  $x_i$  is the  $i$ -th token of  $x$ .  $M_G^E$  denotes the encoder of the baseline translation model ( $M_G$ )

### 4.2 Input representations

In this section, we describe different features we consider for input sentence  $x$  and for the source-side of in-domain sample  $S_d^s$ .

**NMT encoder representations** contain a wealth of information that can be very relevant to the *adaptability* of the model to the new domain. Therefore, we consider it as an important building block of our predictor model. In our implementation, we take the last encoder layer representation of each token, which are then aggregated through a pooling component in a single vector  $\phi_{enc}(x)$  (*Encoder pooling component* at Figure 2).

**Corpus-level features**  $\xi(S_d^s)$  allow us to characterize the in-domain sample  $S_d^s$  with respect to its size, diversity, and similarity to pretrained data  $G$ . In the simplest case we consider sample size (amount of instances) as a single corpus-level feature. In addition we add some of the features used by Xia et al. (2020); Kolachina et al. (2012) namely: (1) amount of tokens in  $S_d^s$ ; (2) vocabulary overlap ratio between  $G$  and  $S_d^s$ ; (3) the average sentence length (in characters, in tokens); (4) the number of unique tokens in  $S_d^s$ ; and (5) type token ratio (Richards, 1987)<sup>7</sup>.

**Instance difficulty features (DF).** The quality of the translation depends on how difficult input sentence is for the NMT system.

Difficulty features  $\phi_{DF}(x)$  include model-based uncertainty functions from Zhao et al. (2020): (1)

training domains.

<sup>6</sup> $T_d$  does not contain the training sentence for adapted model.

<sup>7</sup>We apply log-scale to all feature values to reduce the variability in feature values.

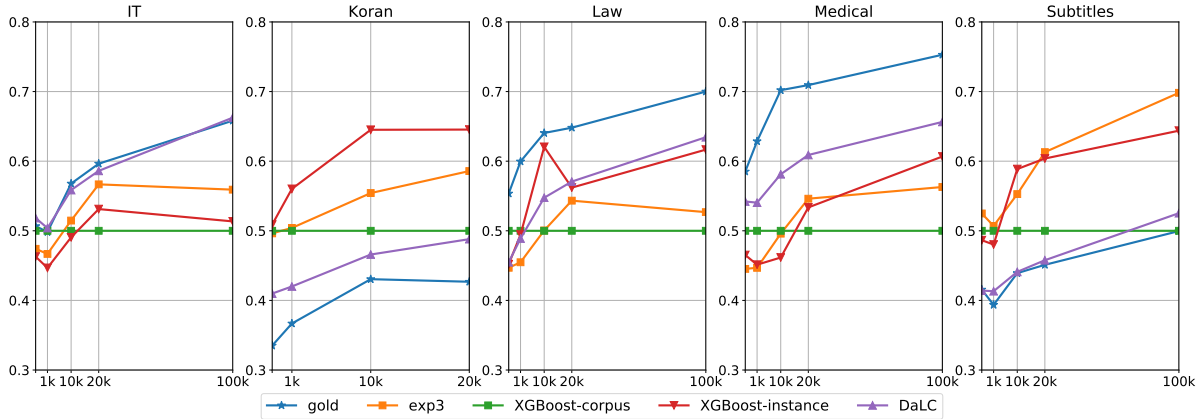


Figure 3: Learning curves provided by DaLC compared against *gold* predictions, baseline predictors (*exp3*, XGboost-corpus, XGboost-instance) for German-English/FT.

least-confidence score, (2) margin score (3) average token entropy. Those features rely on the pretrained model  $M_G$ , and therefore could be seen as redundant with encoder representations. However, as we demonstrate in Section 6 they turn out to be helpful in some cases. Finally, we consider (4) the cross-lingual cosine similarity score between representation of the source sentence  $x$  and its translation  $M_G(x)$ . These representations are obtained from an external pretrained multilingual sentence embedding model, LaBSE (Feng et al., 2020).

### 4.3 Domain Adaptation Learning Curve (DaLC) predictor

DaLC predictor corresponds to the model depicted by Figure 2. It contains two main components: (1) *encoder pooling component* that processes NMT encoder representations (given as a sequence of vectors) and produces a single vector  $\phi_{enc}(x)$ . (2) *fusion layer* combining encoder representations  $\phi_{enc}(x)$  with other pre-computed instance-level features  $\phi_{DF}(x)$  and corpus-level features  $\xi(S_d)$ .

In our experiments, we use the multi-filter CNN architecture proposed by Kim (2014) that is widely used in text classification tasks as *encoder pooling component*. For *fusion component*, we simply stack  $K$  feed-forward layers followed by ReLU and Sigmoid activation functions at the final layer.

## 5 Experimental settings

### 5.1 Data and Evaluation

**English-German.** We rely on data described in Section 3. We use development split (2K sentences) for *predictor training*: it is randomly split into train (80%) and validation (20%) sets.

The predictor is evaluated on test split portion for each domain (2K sentences). The results reported for each domain are obtained with the predictor trained in Leave-one-out settings (e.g., predictor trained on *Law*, *IT*, *Koran*, *Medical* is evaluated on *Subtitles*). Such evaluation allows us to mimic real-life scenario where we need to predict performance for a completely new domain which is not known at training time.

The evaluation of the predictor is done by measuring error between the predicted score and the ground truth score (measured by *mean chrF*) across all the anchor points. Following Kolachina et al. (2012), we report Root Mean Square Error (RMSE) across all the available test anchor points. In addition, we report absolute error at each anchor point (when possible) to allow for finer-grained analysis of the results. Each experiment is repeated 5 times with different random seeds, and an average across all runs is reported.

**English-Korean.** We consider five specialized domains (*technology*, *finance*, *travel*, *sports* and *social science*) publicly available from AI-Hub<sup>8</sup>. The size of the validation and test sets are 10k and 5k, respectively. We randomly sample  $S_{n,d}$  of 2K sentences following the settings of English-German experiment. Detailed information of baseline NMT models and DA models for English-Korean is described in Appendix A. We adopt the same evaluation scenario as in English-German (leave-one-out settings with RMSE evaluation)

<sup>8</sup><https://aihub.or.kr/aihub-data/natural-language/about>

	IT	Koran	Law	Medical	Subtitles	Avg
German-English/FT						
<i>exp3</i>	0.125	0.292	0.305	0.397	0.321	0.288
XGboost-corporus	0.197	0.235	0.305	0.435	0.157	0.266
XGboost-instance	0.084	0.201	0.062	0.176	0.126	0.130
DaLC	<b>0.009</b>	<b>0.058</b>	<b>0.057</b>	0.094	<b>0.015</b>	<b>0.047</b>
DaLC / DF	0.011	0.065	0.058	0.117	0.022	0.055
DaLC / corporus	0.049	0.045	0.097	0.117	0.052	0.072
DaLC / NMTEnc	0.025	0.148	0.085	<b>0.081</b>	0.061	0.080
English-German/FT						
<i>exp3</i>	<b>0.035</b>	0.180	<b>0.116</b>	0.114	0.112	0.112
XGboost-corporus	0.111	<b>0.081</b>	0.169	0.169	<b>0.029</b>	0.112
XGboost-instance	0.072	0.157	0.159	0.116	0.09	0.119
DaLC	0.048	0.107	0.123	0.041	0.057	0.075
DaLC / DF	0.065	0.102	0.123	0.044	0.053	0.077
DaLC / corporus	0.048	0.086	0.126	0.043	0.063	<b>0.073</b>
DaLC / NMTEnc	0.043	0.169	0.125	<b>0.016</b>	0.095	0.090
German-English/Adapter						
<i>exp3</i>	0.055	0.175	0.100	0.169	0.141	0.128
XGboost-corporus	0.079	0.137	0.115	0.166	0.083	0.116
XGboost-instance	0.100	0.145	0.092	0.162	0.097	0.119
DaLC	<b>0.022</b>	<b>0.045</b>	<b>0.080</b>	0.102	<b>0.019</b>	<b>0.054</b>
DaLC / DF	0.024	0.057	0.082	0.109	0.020	0.058
DaLC / corporus	0.030	0.048	0.087	0.109	0.036	0.062
DaLC / NMTEnc	0.079	0.168	0.081	<b>0.068</b>	0.089	0.097

Table 1: RMSE for DaLC predictor compared against *exp3* and XGboost baselines on De-En and En-De directions, where FT and Adapters correspond to the NMT Domain Adaptation method in each experiment.

## 5.2 Baseline predictor models

Traditionally predictor models are evaluated against a naive baseline predicting the *mean* over observations used for training. However, such baseline does not make much sense in the context of learning curve prediction since it is unable to extrapolate to new anchor points.

***exp3* baseline.** *exp3* is a 3-parameter function that is defined by  $y = c - e^{(-ax+b)}$ . Kolachina et al. (2012) has identified this function as a good candidate for SMT learning curve prediction fitting. In our experiment, we fit this function through least-squares algorithm to all the observations we dispose across all the domains and anchor points (19 points). This function can be seen as an extension of *mean* baseline allowing to extrapolate to unobserved anchor points.

**XGboost-based baselines.** Following Xia et al. (2020) we also use gradient boosting trees model (Friedman, 2000), implemented in XGboost (Chen and Guestrin, 2016). *XGboost-corporus* baseline corresponds to the XGboost model trained with the corpus-level features: this baseline is comparable with the one used by Xia et al. (2020). We also compare our model against *XGboost-instance* baseline which corresponds to the XGboost model trained with the full set of features (section 4.2)

that DaLC predictor is trained with. NMT encoder representations are squeezed in a single vector via min-max pooling and fed to XGboost along with other instance-level features. Comparing *XGBoost-corporus* and *XGboost-instance* results allows us to decouple the impact of instance-level representations and impact of the predictor model learnt from these representations. Exact details of XGboost training are reported in Appendix B.

## 5.3 DaLC predictor

In our preliminary experiments, we observed that the capacity of *encoder pooling component* does not have much impact on the overall performance. We believe this is because NMT encoder outputs already provide rich contextualized representations of the input sequence. On the other hand, it is important to give enough capacity to the *Fusion layer* which should mix instance-level representations (including NMT encoder representation) with corpus-level features. In our experiments, *encoder pooling component* is a single layer multi-filter CNN (with 3 filters of size 2, 3, and 4). *Fusion layer* is composed of 4 feed-forward layers of hidden size 512, followed by ReLU activation, and final feed-forward layer followed by Sigmoid activation that brings the final prediction at 0-1 scale. We use Mean Squared Loss (MSE) for training. We

	Finance	Social	Sports	Tech	Travel	Avg
	Korean-English/FT					
<i>exp3</i>	0.055	0.028	0.018	<b>0.080</b>	0.069	0.050
XGboost-corpus	0.246	0.187	0.205	0.307	0.152	0.219
XGboost-instance	0.035	0.066	0.031	0.092	0.027	0.050
DaLC	0.028	<b>0.007</b>	<b>0.016</b>	0.085	<b>0.010</b>	<b>0.029</b>
DaLC / DF	0.046	0.010	0.020	0.086	<b>0.010</b>	0.034
DaLC / corpus	0.037	0.009	0.019	0.104	0.017	0.037
DaLC / NMTEnc	<b>0.025</b>	0.046	0.021	0.089	0.011	0.038

Table 2: RMSE for DaLC predictor compared against *exp3* and XGboost baselines on the finetuning method with the Ko-En direction.

apply early stopping criteria with the patience of 10 epochs. We provide more training hyperparameters in the Appendix E.

## 6 Results

### 6.1 Prediction for observed anchor points

Table 1 reports the results across different domains for English-German and German-English. Table 2 reports the results for Korean-English. We compare corpus-level baselines (*exp3* and *XGboost-corpus*) against different instance-level predictors: (i) our *DaLC predictor* relying on the full set of instance- and corpus-level features described in the section 4.2; (ii) *XGboost-instance* (section 5.2) relying on same features as DaLC; (iii) Ablation of different groups of features from full model (DaLC/DF, DaLC/corpus or DaLC/NMT encoder)

For each domain, we report RMSE when comparing predicted *mean chrF* to the gold *mean chrF*<sup>9</sup> score averaged across all the anchor points (0, 1k, 10k, 20k and 100k).

**Instance-level vs corpus-level.** We note that instance-level models generally outperform corpus-level models for most of domains and language pairs. An exception is English-German direction, where *XGboost-corpus* leads to better prediction than *XGboost-instance* in *Koran* and *Subtitles* domain. According to additional visualisation of these result in Figure 3 and Figure 7 (in Appendix) we see that *XGboost-corpus* model actually fails to learn meaningful patterns as it predicts the same score (0.5) for all the domains across all the anchor

points<sup>10</sup>. It leads to lower RMSE for *Koran* and *Subtitles* domains only because the gold *mean chrF* for those sets is very close to 0.5. Therefore, even if the instance-level leads to higher RMSE it provides more accurate predictions overall as shown by Avg column in Table 1.

We observe that while DaLC reaches lowest RMSE across all the domains for German-English and Korean-English, it is not necessarily true for English-German. We note however, that DaLC performance varies less across domains, and reaches overall best performance (reflected by Avg column) which means that it is less influenced by overall *mean* performance (as opposed to corpus-level models), and is able to better exploit instance-level representations that favour knowledge transfer across domains.

In addition in the Appendix C we report computation cost for instance-level and corpus level models.

**Impact of different features.** We note that the impact of DF and corpus features varies across the domains. One clear trend is that NMT encoder features seem to be important for predictor quality. An interesting exception is the *Medical* domain, where the removal of NMT encoder representations seem to reach the best RMSE. Furthermore, Appendix J provides an in-depth ablation study for each feature used in the model.

**Impact of DA algorithm.** We examine the effectiveness of the proposed method on the different adaptation algorithm, comparing adapter layers (Bapna and Firat, 2019) to full finetuning. The Adapter layer is a small module inserted on top of each encoder and decoder block and updated only

<sup>9</sup>The gold mean chrF corresponds at anchor point  $K$  for domain  $d$  corresponds to the actual value of *mean chrF* obtained after adaptation of NMT model with  $K$  samples from domain  $d$ .

<sup>10</sup>We believe this might be due to the very small amount of corpus-level training samples.

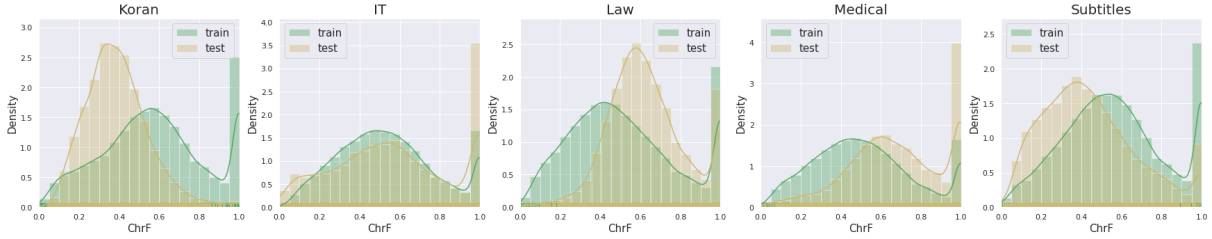


Figure 4: ChrF score distribution of training and test datasets across each experiment case in the German-English FT experiment. Each histogram shows the difference between the training and test ChrF score distributions in each leave one out setting. For instance, the ChrF score distribution of the Koran column obtains where training distribution denotes all training anchor points of IT, Law, Medical and Subtitles domain, and test distribution represent chrF score in all anchor points of the Koran domain.

with in-domain samples (while keeping the rest of the model frozen). The hidden dimension size of the adapter is 1024. Details of the adapted models is provided in Table 5.

Table 1 demonstrates the quality of predictions domain adapted model via Adapter layers in the De-En direction. Comparing to FT results on the De-En direction, we notice that the predictor quality behaves similarly for both DA methods. It indicates that DaLC can be extended to other DA methods.

**Impact of other factors.** We report additional experiments such as depending on the number of domains for training predictor and performance on a mixing of two different domains in Appendix H and Appendix D.

## 6.2 Interpolation and Extrapolation of DA performance

One of practical and possible scenarios is prediction of DA performance for the anchor points that were not observed in the training data. Table 3 shows the accuracy of our predictor for interpolation<sup>11</sup> ( $3k$  and  $40k$ ) and extrapolation<sup>12</sup> ( $160k$ ) scenarios for *Subtitles* test domain in the De-En direction. We recall that the predictors have been trained on  $0k$ ,  $10k$ ,  $20k$  and  $100k$  anchor points for *IT*, *Medical*, *Law* and *Koran* domains.

We report the absolute error with respect to gold *mean chrF* (MAE) for these specific anchor points. We can see that DaLC achieves significantly lower error compared to other baselines. We note however that extrapolation error ( $160k$ ) is higher compared to interpolation errors ( $3k$  and  $40k$ ).

<sup>11</sup>Interpolation: prediction for the unseen anchor points that are within the range of observed anchors.

<sup>12</sup>Extrapolation: prediction for the anchor point that lies beyond observed anchors.

model	3K	40k	160k
<i>exp<sub>3</sub></i>	0.1087	0.1640	0.1934
XGboost-instance	0.1546	0.0922	0.1125
DaLC	<b>0.0063</b>	<b>0.0080</b>	<b>0.0413</b>

Table 3: Results for interpolation and extrapolation of learnt models to new (unseen in training) anchor point: we report the absolute value of the difference between gold and predicted values (subtitles domain)

## 7 Analysis and Discussion

The results reported in previous section suggest that overall DaLC prediction error depends a lot on the nature of the test domain (as seen on Figures 3, 7 and Table 1). Thus, all the predictors tend to overestimate the *mean chrF* score on *Koran* domain, or underestimate the score on *Law* domain. In this section we try to analyse this phenomenon and explain such behaviour.

### 7.1 Train/test data distribution

Figure 4 provides visualization of training and test distribution for each domain. We can see that in the case of *Koran* and *Law* domains there is a highest shift between train and test distribution of chrF scores: *Koran* has more low-quality translations (low chrF value) compared to its training domains, while for *Law* it is the opposite. This discrepancy leads to underestimated scores for *Law* domain and overestimated scores for *Koran*. On the other hand, we can see the predictor has higher accuracy in other domains that have a similar distribution between training and test distributions. We observe the same patterns for the En-De direction (Figure 6), or when extending our framework to more training domains (Appendix H).



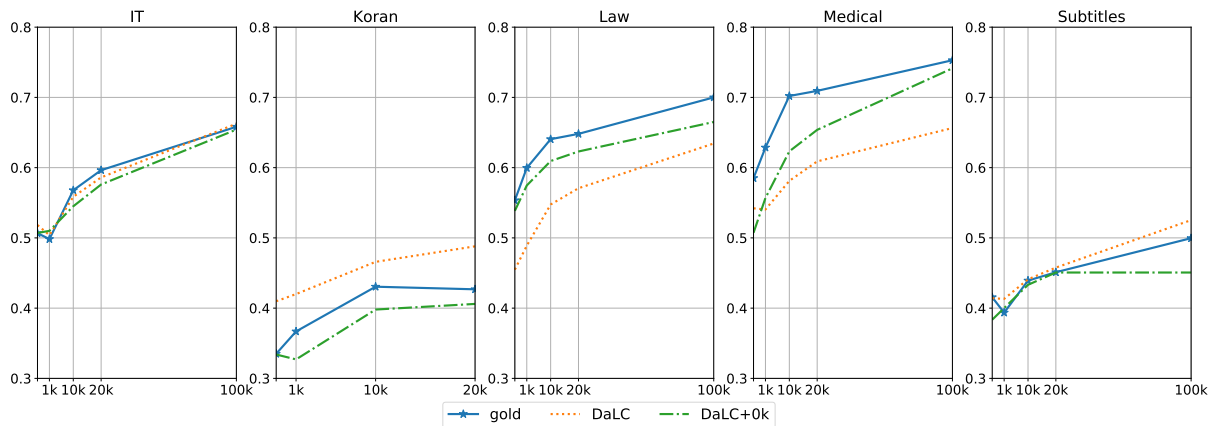


Figure 5: Comparison of learning curve prediction with DaLC and DaLC +with 0k on the De-En direction.

## 7.2 Adding 0 anchor point

The above finding regarding train/test discrepancy leading to a higher prediction error implies that even though our model is able to exploit instance-level representations to some extent (it achieves lower prediction error compared to corpus-level baseline), it is still heavily impacted by overall data distribution. One possible explanation for such behaviour would be the fact that source-level representations may not contain enough information for model to rely on to predict future translation quality. A simple example would be a source sentence that should be either translated in “formal” or “informal” manner depending on what the target domain is. When the model lacks information about target language distribution for the new domain it simply has no mean to learn faithful predictor.

In this experiment, we consider the scenario where we have access to a small sample of parallel data (2K sentences) for the test domain. When such sample is available the predictor training data can be enriched with 0-anchor point samples corresponding to the translations produced by our baseline model and their corresponding chrF scores.

Figure 5 and Figure 8 in Appendix demonstrate the results of this experiment (Table 13 reports RMSE scores for the reader interested in more in-depth analysis). We see that addition of 0k anchor point significantly improves the learning curves obtained by DaLC predictor. This confirms our hypothesis that relying on monolingual in-domain sample may limit the predictors’ performance for certain domains. Adding a small parallel sample to obtain 0k anchor point instances seems to be effective work around for this problem.

## 8 Conclusion

In this work, we formulate a problem of Domain Adaptation Learning Curves prediction for NMT as instance-level learning framework. We demonstrate that it is possible to learn reasonable learning curve prediction model with a very small amount of NMT model instances via instance-level learning rather than corpus-level learning that most previous works rely on. We propose a DaLC model relying on NMT encoder representations, combined with various instance and corpus-level features. We show that such model is able to achieve good results with small amount of pretrained model instances. We perform in-depth analysis of the results for the domains where the predictor was less successful and conclude that the capacity of the predictor relying on the source-side sample only can be limited for some domains. Further analysis of characteristics of such domains could lead to better Domain Adaptation strategies. Finally, we believe it will be interesting to deepen the connections between Active Learning framework and Learning Curves prediction frameworks which could mutually help one another.

## 9 Acknowledgement

The authors would like to thank Seunghyun S. Lim, Hyunjoong Kim and Stéphane Clinchant for valuable comments.

## References

Roei Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 7747–7763.

- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation for neural machine translation](#). In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1538–1548.
- Loïc Barrault, Magdalena Biesialska, Ondřej Bojar, Marta R. Costa-jussà, Christian Federmann, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Matthias Huck, Eric Joanis, Tom Kocmi, Philipp Koehn, Chi-kiu Lo, Nikola Ljubešić, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Santanu Pal, Matt Post, and Marcos Zampieri. 2020. [Findings of the 2020 conference on machine translation \(WMT20\)](#). In *Proc. the Conference on Machine Translation (WMT)*, pages 1–55.
- Alexandre Berard, Ioan Calapodescu, and Claude Roux. 2019. [Naver Labs Europe’s systems for the WMT19 machine translation robustness task](#). In *Proc. the Conference on Machine Translation (WMT)*, pages 526–532.
- Tianqi Chen and Carlos Guestrin. 2016. [Xgboost: A scalable tree boosting system](#). In *Proc. ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, page 785–794, New York, NY, USA. Association for Computing Machinery.
- Maksym Del, Elizaveta Korotkova, and Mark Fishel. 2021. [Translation transformers rediscover inherent data domains](#). In *Proc. the Conference on Machine Translation (WMT)*, pages 599–613.
- Zi-Yi Dou, Antonios Anastasopoulos, and Graham Neubig. 2020. [Dynamic data selection and weighting for iterative back-translation](#). In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5894–5904. Proc. the Annual Meeting of the Association for Computational Linguistics (ACL).
- Zied Elloumi, Laurent Besacier, Olivier Galibert, Juliette Kahn, and Benjamin Lecouteux. 2018. [Asr performance prediction on unseen broadcast programs using convolutional neural networks](#). In *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5894–5898.
- Hady Elsahar and Matthias Gallé. 2019. [To annotate or not? predicting performance drop under domain shift](#). In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2163–2173.
- Fangxiaoyu Feng, Yinfei Yang, Daniel Cer, Naveen Arivazhagan, and Wei Wang. 2020. [Language-agnostic bert sentence embedding](#). *arXiv preprint arXiv:2007.01852*.
- Markus Freitag and Yaser Al-Onaizan. 2016. [Fast domain adaptation for neural machine translation](#). *arXiv preprint arXiv:1612.06897*.
- Jerome H. Friedman. 2000. [Greedy function approximation: A gradient boosting machine](#). *Proc. the Annals of Statistics*, 29:1189–1232.
- Behrooz Ghorbani, Orhan Firat, Markus Freitag, Ankur Bapna, Maxim Krikun, Xavier Garcia, Ciprian Chelba, and Colin Cherry. 2021. [Scaling laws for neural machine translation](#). *arXiv preprint arXiv:2109.07740*.
- Mitchell A Gordon, Kevin Duh, and Jared Kaplan. 2021. [Data and parameter scaling laws for neural machine translation](#). In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5915–5922.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *arXiv preprint arXiv:2001.08361*.
- Yoon Kim. 2014. [Convolutional neural networks for sentence classification](#). In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751.
- Prasanth Kolachina, Nicola Cancedda, Marc Dymetman, and Sriram Venkatapathy. 2012. [Prediction of learning curves in machine translation](#). In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 22–30.
- Minh-Thang Luong and Christopher D Manning. 2015. [Stanford neural machine translation systems for spoken language domains](#). *Proc. the International Conference on Spoken Language Translation (IWSLT)*.
- Nitika Mathur, Johnny Wei, Markus Freitag, Qingsong Ma, and Ondřej Bojar. 2020. [Results of the wmt20 metrics shared task](#). In *Proc. the Conference on Machine Translation (WMT)*, pages 688–725.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proc. of The Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Weijing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 311–318.
- Álvaro Peris and Francisco Casacuberta. 2018. [Active learning for interactive neural machine translation of data streams](#). In *Proc. the Conference on Computational Natural Language Learning (CoNLL)*, pages 151–160.
- Maja Popović. 2015. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proc. the Workshop on Statistical Machine Translation (WMT)*, pages 392–395.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proc. the Conference on Machine Translation (WMT)*, pages 186–191.

- Brian Richards. 1987. [Type/token ratios: What do they really tell us?](#) *Journal of child language*, 14(2):201–209.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units.](#) In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1715–1725.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS.](#) In *Proc. the International Conference on Language Resources and Evaluation (LREC)*, pages 2214–2218.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need.](#) In *Proc. the Advances in Neural Information Processing Systems (NeurIPS)*.
- Wei Wang, Ye Tian, Jiquan Ngiam, Yinfei Yang, Isaac Caswell, and Zarana Parekh. 2020. [Learning a multi-domain curriculum for neural machine translation.](#) In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 7711–7723.
- Mengzhou Xia, Antonios Anastasopoulos, Ruochen Xu, Yiming Yang, and Graham Neubig. 2020. [Predicting performance for natural language processing tasks.](#) In *Proc. the Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 8625–8646.
- Zihuiwen Ye, Pengfei Liu, Jinlan Fu, and Graham Neubig. 2021. [Towards more fine-grained and reliable nlp performance prediction.](#) In *Proc. the Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, pages 3703–3714.
- Runzhe Zhan, Xuebo Liu, Derek F. Wong, and Lidia S. Chao. 2021. [Meta-curriculum learning for domain adaptation in neural machine translation.](#) volume 35, pages 14310–14318.
- Ye Zhang, Matthew Lease, and Byron Wallace. 2017. [Active discriminative text representation learning.](#) In *Proc. the AAAI Conference on Artificial Intelligence (AAAI)*, volume 31.
- Yuekai Zhao, Haoran Zhang, Shuchang Zhou, and Zhihua Zhang. 2020. [Active learning approaches to enhancing neural machine translation.](#) In *Proc. of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1796–1806.

## A NMT model training and domain adaptation hyper parameters

### A.1 Baseline

**English-German** Our baseline models rely on `transformer-base` architecture (Vaswani et al., 2017) trained on WMT-20 German-English dataset. For English-German direction we used inline casing (Berard et al., 2019) and tokenized text with BPE (Sennrich et al., 2016) using 24K vocabulary size encoding joint vocabulary shared between two languages. We also share embedding parameters between source and target. We use Fairseq toolkit (Ott et al., 2019) for training the baseline model with default training parameters for `transformer-base` architecture as proposed by Vaswani et al. (2017). Specifically,  $M_G$  consists of six layers, 512 units and eight heads. We utilize the same Adam warm-up optimizer as in the original paper. Note that the baseline model for Finetuning and Adapter experiments in the same language direction are identical regardless of Domain Adaptation methods.

**Korean-English** We follow the same experiment setting described in Appendix. E. We utilize `transformer-big` architecture (Vaswani et al., 2017) trained on the Ko-En direction. Training datasets of baseline model is constructed by using six categories datasets<sup>13</sup> in AI-Hub<sup>14</sup>. The total number of training dataset is 1.51M parallel sentences. We set the size of BPE as 32K.

### A.2 Domain Adaptation

Hyper-parameters used for Domain Adaptation models are reported in the Table 4. We use early stopping criteria on the validation loss performance in order to avoid overfitting to the data. Note that at this stage we used greedy decoding to compute chrF scores. We verified that the actual learning curves with greedy decoding behaves similarly to the learning curves with beam search. We stick to greedy decoding because we believe that beam search brings additional complexity to the performance prediction. We leave the impact of decoding method on the performance prediction for future work.

Table 5 and Table 8 report the BLEU/mean chrF scores reached by each of our models on each lan-

<sup>13</sup>We aggregate news, dialogue, colloquial, Korean culture, ordinance and website.

<sup>14</sup><https://aihub.or.kr/aidata/87>

	En-De	De-En	Ko-En
Learning rate	0.001	0.001	0.001
batch size	4k	2k	2k

Table 4: Parameters used for Domain Adaptation of NMT system in each direction, where the batch size denotes the tokens per batch. En-De/FT and En-De/Adapter utilize the same learning rate and the batch size.

guage. We note that the model adapted via finetuning follows the same trend as the one adapted via adapter layers. Table 6 reports best achieved loss for each domain and each anchor point, as well as the number of epochs required to reach it. One needs to keep in mind that the size of a single epoch varies across anchor points (it is 10 times bigger for 10k anchor point than for 1k anchor point). Note that loss is directly comparable between different domains and different anchor points since it corresponds to Cross-entropy loss relying on the same vocabulary.

## B Implementation details of XGboost

We run the XG-boost regressor using XGBoost library<sup>15</sup>. The learning rate of the XG-boost is set to 0.1. We follow the default regularization to alleviate the overfitting problem. The objective function of the XG-boost is RMSE. Regarding the model parameters, the number of trees and the maximum depth of trees are 100 and 10, respectively.

## C Computation time and memory usage

We compare the computation time and memory usage of DaLC with other models on the learning curve prediction task in Table 7. Although, corpus-level frameworks(i.e., XGboost-corpus and exp3) require less computation cost and memory resources, the performances of these models with small number of anchor points are very far from the real value described in Table 1. The most computational heavy part of our model is the computation of anchor points used to train the predictor model. Anchor point implies training a Domain Adapted model with in-domain samples (200 300 sec in best case scenario of 1000 in-domain samples and up to 4h for larger in-domain samples). For this reason, instance-level frameworks are more practical in real-life scenarios as they achieve good performance with small amount of anchor points. Note,

<sup>15</sup><https://github.com/dmlc/xgboost>

	IT		Koran		Law		Medical		Subtitles	
	BLEU	chrF	BLEU	chrF	BLEU	chrF	BLEU	chrF	BLEU	chrF
	German-English/FT									
Baseline	29.2	0.506	10.2	0.335	27.8	0.554	32.7	0.585	20.9	0.415
1K	27.8	0.498	13.3	0.367	34.6	0.600	36.6	0.629	21.2	0.393
10K	30.8	0.568	20.2	0.430	40.1	0.641	45.0	0.702	24.9	0.439
20K	33.5	0.596	19.7	0.426	42.6	0.648	47.7	0.709	26.0	0.451
100K	39.5	0.658	-	-	49.0	0.696	55.1	0.753	30.8	0.500
	English-German/FT									
Baseline	32.1	0.544	11.9	0.358	36.8	0.627	34.5	0.599	21.8	0.472
1K	33.4	0.576	16.3	0.393	39.2	0.643	37.7	0.634	23.9	0.472
10K	34.6	0.605	21.5	0.441	42.3	0.668	41.9	0.664	25.2	0.482
20K	35.9	0.619	23.4	0.452	43.3	0.675	42.6	0.668	25.2	0.484
100K	38.2	0.653	-	-	46.7	0.702	46.6	0.694	25.6	0.490
	German-English/Adapters									
Baseline	29.2	0.506	10.2	0.335	27.8	0.554	32.7	0.585	20.9	0.415
1K	30.0	0.520	12.5	0.341	33.1	0.562	33.0	0.613	16.7	0.374
10K	30.2	0.562	17.2	0.387	38.6	0.614	41.9	0.670	21.5	0.416
20K	31.3	0.585	19.0	0.401	40.0	0.629	44.2	0.691	23.5	0.445
100K	35.2	0.638	-	-	46.0	0.673	51.1	0.730	25.4	0.476

Table 5: Domain Adaptation performance across domains for English-German and German-English models as measure either in BLEU or in chrF scores.

	IT		Koran		Law		Medical		Subtitles	
	Epoch	loss	Epoch	loss	Epoch	loss	epoch	loss	epoch	loss
	English-German/FT									
1K	4	3.718	12	4.228	4	3.033	5	3.203	3	3.840
10K	12	3.502	20	3.680	10	2.929	12	3.034	9	3.784
20K	11	3.433	20	3.489	9	2.881	10	2.971	7	3.762
100K	11	3.242	-	-	13	2.759	12	2.809	6	3.691

Table 6: Domain Adaptation convergence statistics. We report the number of epochs when the Domain Adaptation reached the best validation loss, as well as the value of best validation loss. Note that each epoch corresponds to different amount of updates for different anchor point due to the difference in the in-domain samples used for DA. The loss however is comparable across different anchor points and different domains since it is always based on the same vocabulary.

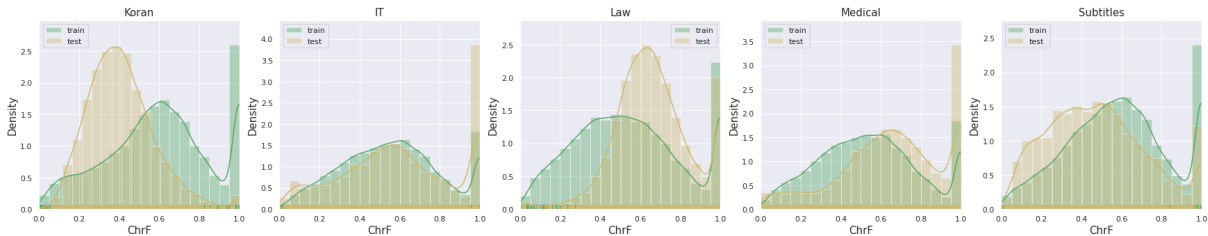


Figure 6: ChrF score distribution of training and test datasets across each experiment case in the English-German FT experiment. Each histogram shows the difference between the training and test ChrF score distributions in each least one out setting. For instance, the training distribution of Koran column obtains from all training anchor points of IT, Law, Medical and Subtitles domain. Test distribution of Koran column represents chrF score in all anchor points of the Koran domain.

Model	Train time(s/epoch)	Test time(s/anchor)	params
exp3	0.001	0.0002	3
XGboost-corpora	0.23	0.0013	100
XGboost-instance	1.44	0.0802	1.03M
DaLC	8.23	0.0914	3.9M

Table 7: Computation time and memory usage of each model, where the test time is the time to inference one anchor point.

that previous works Xia et al. (2020); Kolachina et al. (2012) relied on 50-90 of MT model instances to train reliable corpus-level predictor.

Considering the performance and total training time to obtain training anchor points, instance-level frameworks are feasible solutions to estimate the future performance without finetuning. In instance-level frameworks, DaLC is slightly slower than XGboost-instance but DaLC outperforms other approaches as shown in Table 1. Moreover, DaLC is a very light in terms of parameter size compared to transformer base (65M). Our model can accurately forecast the finetuning performance with only 4% of the total number parameters of transformer base without the need of computation resources for finetuning.

## D Performance on the Mixed domain

We evaluate our model under a loosely defined domain, such as mixing of two domains. To simulate this scenario, we defines a mixed domain by aggregating Tech and TED domain. The samples  $S_{n,d}$  of size  $n$  in the mixed domain uniformly sampled from each domain. The test set of a mixed domain also is constructed with all test sentences of both domains. We train DaLC with anchor points of Finance, Social, Sports and Travel (i.e., Except ‘Tech’ domain) and then evaluate the performance on the mixed domain. Table 10 summarizes the performance on each anchor point in the mixed domain experiment.

Table 9 demonstrates the RMSE errors across anchor points on the defined mixed domain and each domain. As previously, DaLC reaches lower (and more stable) prediction error compared to corpus-level models on averages. Thus, even the corpus-level can reach lower error in some cases (E.g., XGboost-corpora on TED) they still tend to be very close to predicting the mean score of the training data, and therefore their predictions are not stable across domains, while DaLC’s predictions are more stable.

## E Predictor implementation details

The overall architecture of the predictor is described in Figure 2. Our prediction is composed of two parts, the encoder pooling component and the fusion layer component. The encoder pooling component is based on the multi-filter CNN architecture proposed by Kim (2014) which is widely used in text classification tasks. It encodes the sentence to latent feature  $h_i$ , which contains context information to predict performance.

At first, the model obtains the encoder representation  $e_i \in \mathcal{R}^d$  of the  $i$ -th token of the input  $x = (x_1, x_2, \dots, x_n)$  from the encoder of the baseline model ( $M^G$ ), where  $d$  and  $n$  is the size of dimension and the length of  $x$ . The size of hidden dimension  $d$  is 512 in our implementation. Next, we apply the multiple convolution filters across the sequence in the same manner as the original paper, where the window sizes of convolution filters are 2,3 and 4. We then obtain the output  $h_i \in \mathcal{R}^d$  from the max pooling operation along the sequence of convolution outputs. To fuse all features, we concatenate the features of encoder representations  $h_i$ , the pre-computed domain-difficulty features  $\phi_{DL}(x)$  and the corpus-level features  $\xi(S_d)$ . The fusion layer forecasts the performance based on the concatenated features, where the fusion layer is constructed with a five layer feed-forward neural network with non-linear activation functions. We utilize a non-linear function as ReLU, but the last layer of the feed-forward neural network utilize the Sigmoid function to change the output from zero to one. Learning is done via MSE loss between the predicted score and the ground-truth score.

**Training hyperparameters** To optimize the model, we use the learning rate as 0.001 with Adam optimizer. Moreover, the learning rate gradually decreases based on an exponential decay scheduler. All experiments stop the training based on the early stopping with patience as 10. Note that the parameters of  $M^G$  encoder are fixed while training the predictor.

## F Overview of predictor training dataset creation

- We start from the deduplicated train/dev/test splits.
- we sampled 1k, 10k, 20k or 100k samples from the train split

	Finace		Social		Sports		Tech		Travel	
	BLEU	chrF	BLEU	chrF	BLEU	chrF	BLEU	chrF	BLEU	chrF
	Korean-English/FT									
Baseline	49.5	0.718	37.9	0.644	46.2	0.678	27.6	0.583	37.2	0.623
1K	52.5	0.736	41.9	0.674	48.8	0.695	50.9	0.728	39.9	0.645
10K	54.6	0.749	44.5	0.691	51.2	0.712	56.9	0.768	41.4	0.655
20K	55.5	0.754	45.6	0.699	52.2	0.717	58.9	0.780	42.1	0.660
100K	58.0	0.768	48.5	0.718	54.7	0.732	63.2	0.807	43.8	0.672

Table 8: Domain Adaptation performance across domains for Korean-English model as measure either in BLEU or in chrF scores.

	Mixed	TED	Tech
<i>exp3</i>	0.231	0.297	0.136
XGboost-Corpus	0.170	0.070	0.229
XGboost-Instance	0.169	0.196	0.141
DaLC	0.105	0.112	0.099

Table 9: RMSE on the mixed domain experiments, where ‘Mixed’ column means RMSE error on performance estimation in the mixed domain test sentences. ‘TED’ and ‘Tech’ column represents the performance on only test sentences on the corresponding domain in the mixing domain.

	Mixed	TED	Tech
	Korean-English/FT		
Baseline	0.503	0.423	0.583
1K	0.562	0.427	0.697
10K	0.590	0.431	0.748
20K	0.596	0.429	0.763
100K	0.618	0.439	0.795

Table 10: ChrF scores on each test set with a mixed Domain Adaptation case, where the models trained with uniformly sampled TED and Tech training dataset. Note that a Domain Adaptation model on each anchor point(i.e.,row-wise) is the same adaptions model. ‘Mixed’ represents average chrF scores on test sentences of both TED and Tech domains.

- we finetuned English-German baseline NMT system on each of those samples
- We then were able to compute instance-level chrF for each domain, at each anchor point for dev and test splits.

## G Additional results

Table 11 reports prediction errors (RMSE/MAE) across all available anchor points for all domains for **DaLC**. Additionally, we provide results of the learning curves depending on the different combination components of DaLC in Figure 9. Table 12 reports the deviation from the gold predictions for each anchor point, each domain. Figure 7 is the learning curve prediction result in the English-German direction. We also report Table 13 to demonstrate the detail RMSE errors with + with 0k experiments described in Section 7.2

## H Impact of the amount of domains

Elsahar and Gallé (2019) have shown one can obtain better precision at performance drop prediction due to domain shift when the amount of training domains increases. Inspired by this finding, we extend the set of our training domains with additional domains (English-German language pair only): *Bible*, *Tatoeba*, *Medline*, *TAUS*, *PatTR*, and *MuchMore*. We provide some details of these various datasets in the Appendix I.

For each of these datasets, we created a predictor training dataset following the same settings as previously (whole procedure is outlined in the Appendix F).

We train DaLC predictor keeping the same test domains as previously, but extend corresponding training domains with newly introduced domains. We consider following extensions of the training domains: (*Bible*, *Medline*, *PatTR*), (*TAUS*, *Tatoeba*,

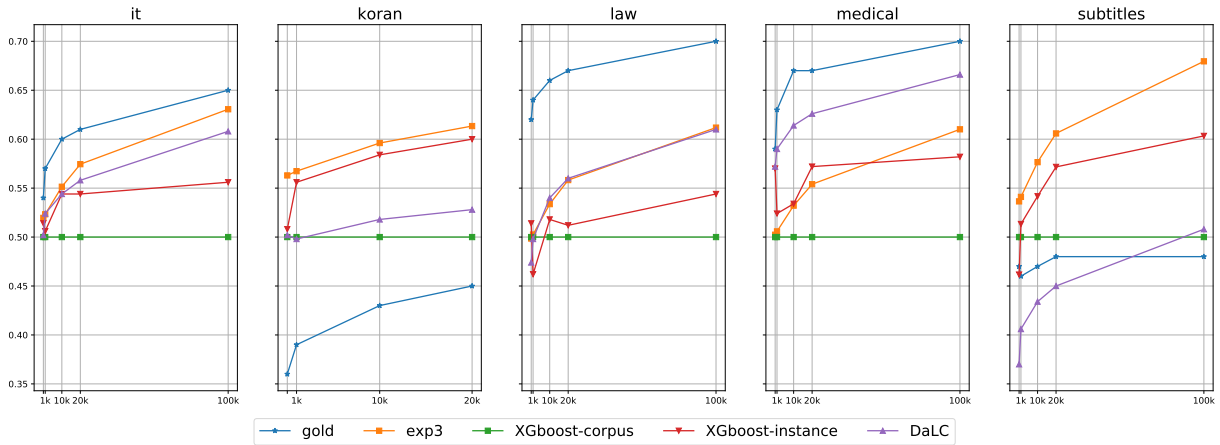


Figure 7: Learning curves provided by DaLC compared against *gold* predictions, baseline predictors (*exp3*, XGboost-corpus, XGboost-instanc) for English-German.

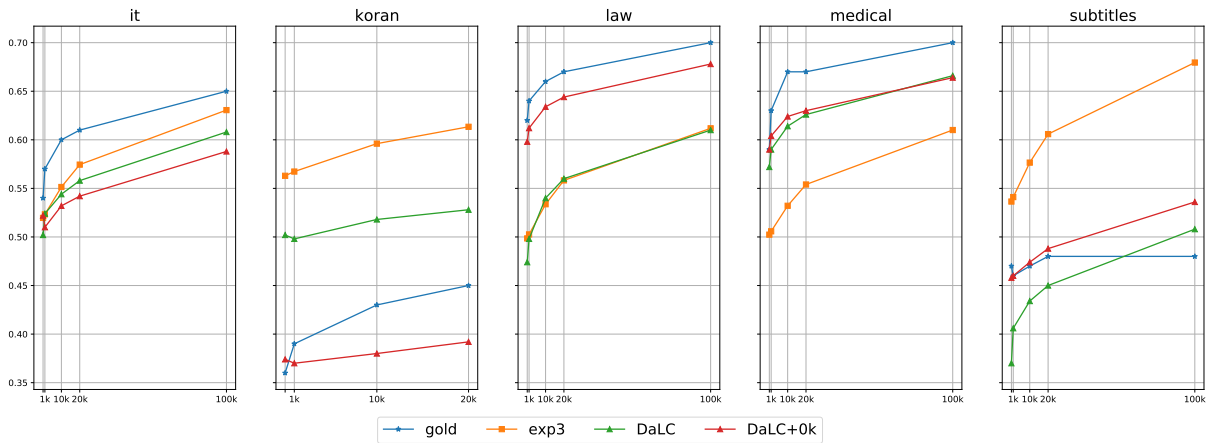


Figure 8: Comparison of learning curve prediction with DaLC and DaLC +with 0k on the English to German

*MuchMore*) or *All* (6 new domains). The goal of this split is to decouple the differences in performance that are due to the amount of training data from those due to the nature of new domains added. Table 15 reports the results. We can see that additional domains indeed improve the overall error. We note that error decreases *Law* domain, but error increases on *Koran* domain. When looking at the result we see that the addition of new domains simply leads to higher predictions scores overall (which is due to overall high chrF scores in the additional domains).

### I Additional datasets

- Bible, Tatoeba from Opus Website<sup>16</sup>
- Medline: consists of abstracts from scientific publications, distributed as part of WMT-20

<sup>16</sup><https://opus.nlpl.eu/>

Biomedical translation challenge<sup>17</sup>

- TAUS<sup>18</sup>: TAUS Corona Crisis Corpora that consists of crawled documents related to Covid-19 crisis.
- PatTR<sup>19</sup>: patents related to medical domain
- MuchMore<sup>20</sup>: scientific medical abstracts obtained from the Springer Link web site.

### J In-depth ablation study

Table 14 reports the results for in-depth ablation study on Subtitles domain. Table 14 shows the

<sup>17</sup><http://www.statmt.org/wmt20/biomedical-translation-task.html>

<sup>18</sup><https://md.taus.net/corona>

<sup>19</sup>[https://ufal.mff.cuni.cz/ufal\\_medical\\_corpus](https://ufal.mff.cuni.cz/ufal_medical_corpus)

<sup>20</sup>[https://ufal.mff.cuni.cz/ufal\\_medical\\_corpus](https://ufal.mff.cuni.cz/ufal_medical_corpus)



performance of the model after removing the corresponding feature from the input in the DaLC. The result shows that all features contribute to improve the performance in our task and empirically demonstrates the importance of each feature. The normations correspond to

- MS : margin score,
- LC: least-confidence score,
- ATE : average token entropy,
- Labse: LaBSE based cosine similarity score,
- c-length : the average sentence length in characters,
- n-token: the amount of tokens in  $S_d^s$ ,
- l-token : the average sentence length in tokens,
- TTR : type token ratio in  $S_d^s$ ,
- overlap: vocabulary overlap ratio between  $G$  and  $S_d^s$ ,
- n-vocab :the number of unique tokens in  $S_d^s$ ,

where the each term is mentioned at the Section 4.2.

domain	Mean		$exp_3$ -fit		XGboost-instance		DaLC/DF		DaLC	
	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE
koran	0.328	0.164	0.292	0.145	0.201	0.200	0.064	0.065	<b>0.058</b>	<b>0.056</b>
it	0.073	0.028	0.125	0.049	0.084	0.076	0.011	0.009	<b>0.009</b>	<b>0.009</b>
medical	0.366	0.162	0.397	0.176	0.176	0.172	0.117	0.115	<b>0.094</b>	<b>0.082</b>
law	0.248	0.108	0.305	0.134	0.085	0.079	0.058	<b>0.047</b>	<b>0.057</b>	0.053
subtitles	0.308	0.094	0.321	0.139	0.126	0.120	0.022	0.019	<b>0.015</b>	<b>0.007</b>

Table 11: Global Learning Curve prediction error: we report RMSE/MAE across all the anchor points available for each domain. De-En translation, Adaptation method: finetuning

	it	koran	law	medical	subtitles
0	-0.011 (0.506)	-0.075 (0.335)	0.082 (0.554)	0.043 (0.585)	0.001 (0.415)
1000	-0.005 (0.498)	-0.053 (0.367)	0.072 (0.600)	0.088 (0.629)	-0.020 (0.393)
10000	0.010 (0.568)	-0.036 (0.430)	0.039 (0.641)	0.120 (0.702)	-0.002 (0.439)
20000	0.002 (0.596)	-0.061 (0.426)	0.029 (0.648)	0.100 (0.709)	0.007 (0.451)
100000	-0.004 (0.658)	-	0.044 (0.696)	0.096 (0.753)	-0.026 (0.500)

Table 12: Learning Curves prediction results. En-De translation, Adaptation method: finetuning Each cell reports gold *mean chrF* score and prediction model deviation from gold mean chrF ( prediction error ).

	German-English/FT					English-German/FT				
	IT	Koran	Law	Medical	Subtitles	IT	Koran	Law	Medical	Subtitles
XGboost-instance	0.092	0.209	0.080	0.176	0.151	0.075	0.158	0.155	0.105	0.083
DaLC	<b>0.006</b>	0.051	0.050	0.102	<b>0.017</b>	0.064	0.111	0.128	0.032	0.050
Adding 0k anchor point of the test domain										
XGboost + with 0k	0.900	0.144	0.073	0.156	0.900	0.076	0.152	0.160	0.108	0.086
DaLC + with 0k	0.007	<b>0.032</b>	<b>0.029</b>	<b>0.060</b>	0.024	<b>0.059</b>	<b>0.028</b>	<b>0.023</b>	<b>0.016</b>	<b>0.035</b>

Table 13: RMSE for DaLC predictor and XGboost baseline when adding 0k anchor point of a test domain as additional training data on English-German language pair. +with 0k represents the model utilizes the 0k anchor point of the given test domain in training. We report RMSE across all anchor points of the test domain excluding the 0k anchor point.

	Full	$\phi_{DF}$				$\xi$					
	DaLC	wo MS	wo LC	wo ATE	wo Labse	wo c-length	wo n-token	wo l-token	wo TTR	wo overlap	wo n-vocab
RMSE	0.0097	0.0175	0.0103	0.0186	0.0112	0.0179	0.0174	0.0241	0.0263	0.0194	0.0119
MAE	0.0085	0.0141	0.0121	0.0149	0.0095	0.0153	0.0147	0.0176	0.0176	0.0159	0.0102

Table 14: Ablation study of DaLC. AL with the subtitles domain, where the model train with the other domains case, such as koran, medical, it and law domains. ‘wo’ indicates removal of corresponding feature from the input of the model. For example, ‘wo lf’ indicates removing the least confidence score in the input.

	IT	Koran	Law	Medical	Subtitles	Avg
4 domains	0.054	0.114	0.124	0.032	0.06	0.077
+ Bible + Medline + PatTR	0.033	0.128	0.076	0.031	0.044	0.062
+ TAUS + Tatoeba + MuchMore	0.016	0.124	0.083	0.012	0.092	0.065
+ All	0.028	0.136	0.061	0.018	0.089	0.066

Table 15: Impact of the number of domains on the predictors performance (En-De, FT)

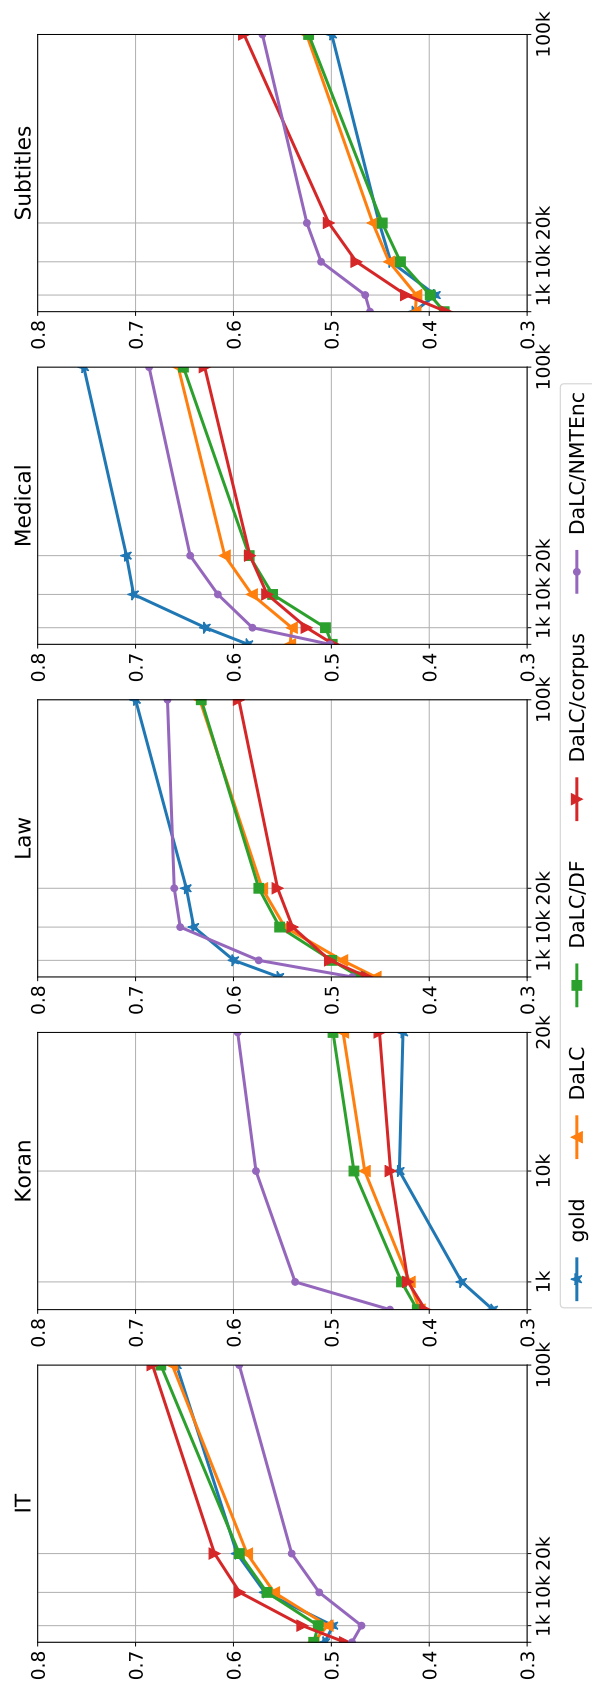


Figure 9: Learning curves for gold scores and predictors combining different components (please refer to section 6 for details of those combinations). Best viewed in color.