

Data Selection Curriculum for Neural Machine Translation

Tasnim Mohiuddin^{1,2*} Philipp Koehn² Vishrav Chaudhary²
James Cross² Shruti Bhosale² Shafiq Joty¹
Nanyang Technological University¹ Meta AI²
{mohi0004, srjoty}@ntu.edu.sg
{pkoehn, vishrav, jcross, shru}@fb.com

Abstract

Neural Machine Translation (NMT) models are typically trained on heterogeneous data that are concatenated and randomly shuffled. However, not all of the training data are equally useful to the model. Curriculum training aims to present the data to the NMT models in a meaningful order. In this work, we introduce a two-stage training framework for NMT where we fine-tune a base NMT model on subsets of data, selected by both deterministic scoring using pre-trained methods and online scoring that considers prediction scores of the emerging NMT model. Through comprehensive experiments on six language pairs comprising low- and high-resource languages from WMT’21, we have shown that our curriculum strategies consistently demonstrate better quality (up to +2.2 BLEU improvement) and faster convergence (approximately 50% fewer updates).

1 Introduction

The notion of a curriculum came from the human learning experience; we learn better and faster when the learnable examples are presented in a meaningful sequence rather than a random order (Newport, 1990). In the case of machine learning, curriculum training hypothesizes presenting the data samples in a meaningful order to machine learners during training such that it imposes structure in the task of learning (Bengio et al., 2009).

In recent years, Neural Machine Translation (NMT) has shown impressive performance in high-resource settings (Hassan et al., 2018; Popel et al., 2020). Typically, training data of the NMT systems are a heterogeneous collection from different domains, sources, topics, styles, and modalities. The quality of the training data also varies a lot, so as their linguistic difficulty levels. The usual practice of training NMT systems is to concatenate all available data into a single pool and randomly sample

training examples. However, not all of them may be useful, some examples may be redundant, and some data might even be noisy and detrimental to the final NMT system performance (Khayrallah and Koehn, 2018). So, NMT systems have the potential to benefit greatly from curriculum training in terms of both speed and quality.

In this work, we propose a *two-stage* training framework for NMT — *model warm-up* and *model fine-tuning*, where we apply the data-selection curriculum in the later stage. We initially train a base NMT model in the warm-up stage on all available data. In the fine-tuning stage, we adapt the base model on selected subsets of the data. The subset selection is performed by considering data quality and/or usefulness at the current state of the model. We explore two sets of data-selection curriculum strategies — *deterministic* and *online*. The deterministic curriculum uses external measures which require pretrained models for selecting the data subset at the beginning of the model fine-tuning stage and continues training on the selected subset. In contrast, the online curriculum dynamically selects a subset of the data for each epoch without requiring any external measure. Specifically, it leverages prediction scores of the emerging NMT model which are the training by-product.

For picking the data subset in the online curriculum, we investigate two approaches of *data-selection window* — static and dynamic. Even though the size of the data-selection window is *constant* throughout the training in the static approach, the samples in the selected subset *vary* from epoch-to-epoch due to the change in their prediction scores. In contrast, we *change* the data-selection window size in the dynamic approach by either expanding or shrinking.

Comprehensive experiments on six language pairs (12 translation directions) comprising low- and high-resource languages from WMT’21 (Akhbardeh et al., 2021) reveal that our curricu-

*Work done while Tasnim was interning at Meta AI

lum strategies consistently demonstrate better performance compared to the baseline trained on all the data (up to +2.2 BLEU). We observe bigger gains in the high-resource pairs compared to the low-resource ones. Interestingly, we find that the online curriculum approaches perform on par with the deterministic approaches while not using any external pretrained models. Our proposed curriculum training approaches not only exhibit better performance but also converge much faster requiring approximately 50% fewer updates.

2 Proposed Framework

Let s and t denote the source and target language respectively, and $\mathcal{D}_g = \{(x_i, y_i)\}_{i=1}^N$ denote the general-domain parallel training data containing N sentence pairs with x_i and y_i coming from s and t languages, respectively. Also, let $\mathcal{D}_d \subseteq \mathcal{D}_g$ be the in-domain parallel training data and \mathcal{M} is an NMT model that can translate sentences from s to t . The overall training objective of the NMT model is to minimize the total loss of the training data:

$$\mathcal{J}(\theta) = \sum_{i=1}^N \mathcal{L}(x_i, y_i, \theta) = \sum_{i=1}^N -\log P_{\theta}(y_i|x_i) \quad (1)$$

where $P_{\theta}(y_i|x_i)$ is the sentence-level translation probability of the target sentence y_i for the source sentence x_i with θ being the parameters of \mathcal{M} .

We propose a *two-stage* training framework where in *model warm-up* stage we train \mathcal{M} on general domain parallel data \mathcal{D}_g for K number of gradient updates; K is generally smaller than the total number of updates \mathcal{M} requires for convergence. Then in *model fine-tuning* stage, we adapt \mathcal{M} on selected subsets of in-domain parallel data \mathcal{D}_d . Based on the intuition: “*not all of the training data are useful or non-redundant, some samples might be irrelevant or even detrimental to the model*”, we hypothesize that there exists a $\mathcal{D}_s \subset \mathcal{D}_d$, fine-tuning on which \mathcal{M} will exhibit improved performance.

Our goal is to design a ranking of the training samples which will eventually help us to extract \mathcal{D}_s from \mathcal{D}_d . For this, we investigate two sets of data-selection curriculum strategies — *deterministic* and *online*. Both strategies require a measure of data quality and/or usefulness at the current state of the model to extract \mathcal{D}_s . While the deterministic curriculum uses external measures that require pretrained models, the online curriculum leverages the prediction scores of the emerging NMT models.

2.1 Deterministic Curriculum

In this strategy, we select a $\mathcal{D}_s \subset \mathcal{D}_d$ initially and do not change it during the *model fine-tuning* stage. We first score each parallel sentence pair $(x_i, y_i) \in \mathcal{D}_d$ using an external bitext scoring method. We experiment with three such scoring methods as described below.

- **LASER** This approach utilizes the Language-Agnostic SEntence Representations (LASER) toolkit (Artetxe and Schwenk, 2019), which gives multilingual sentence representations using an encoder-decoder architecture trained on a parallel corpus. We use the sentence representations to *score the similarity* of a parallel sentence pair using the Cross-Domain Similarity Local Scaling (CSLS) measure (Conneau et al., 2017).

$$Score_{\text{laser}}(x_i, y_i) = \text{CSLS}(\text{LASER}(x_i), \text{LASER}(y_i)) \quad (2)$$

- **Dual Conditional Cross-Entropy (DCCE)** Junczys-Dowmunt (2018) proposed this method, which requires two inverse translation models – one forward model (f) and one backward model (b), trained on the same parallel corpus. It then finds the score of a sentence pair (x_i, y_i) by taking the maximal symmetric agreement of the two models which exploits the conditional cross-entropy (H).

$$Score_{\text{dce}}(x_i, y_i) = |H_f - H_b| + \frac{1}{2}(H_f + H_b) \quad (3)$$

where $H_f = -\log P_{\theta_f}(y_i|x_i)$; $H_b = -\log P_{\theta_b}(x_i|y_i)$

- **Modified Moore-Lewis (MML)** MML ranks the sentence pairs based on domain relevance by calculating *cross-entropy difference* scores (Moore and Lewis, 2010; Axelrod et al., 2011). For this, we need to train four language models (LM): **in-** and **general-domain** LMs in both source and target languages. Then we find the MML score of a parallel sentence pair (x_i, y_i) as follows:

$$Score_{\text{mml}}(x_i, y_i) = (H_{s,\text{in}}(x_i) - H_{s,\text{gen}}(x_i)) + (H_{t,\text{in}}(y_i) - H_{t,\text{gen}}(y_i)) \quad (4)$$

$$\text{where } H_{b,C}(z) = -\log P_{b,C}^{\text{LM}}(z)$$

Here, $b \in \{s, t\}$ refers to the bitext side and $C \in \{\text{in}, \text{gen}\}$ refers to the corpus domain. In our experiments, we use the *newscrawl* data as **in-domain** and *commoncrawl* data combined with *newscrawl* as **general-domain** for training the LMs.

After scoring each parallel sentence pair $(x_i, y_i) \in \mathcal{D}_d$ by any of the above methods, we rank \mathcal{D}_d based on the scores. We then pick the

Algorithm 1 Deterministic Curriculum Strategy

Input : General domain corpus \mathcal{D}_g , in-domain corpus $\mathcal{D}_d \subseteq \mathcal{D}_g$, external pretrained bi-text scorer \mathcal{S}

Output : A trained translation model

1. // model warm-up stage
Train a base model \mathcal{M} on general domain corpus \mathcal{D}_g for K number of updates
 2. // model fine-tuning stage
 - (a) **Score** each $(x_i, y_i) \in \mathcal{D}_d$ using \mathcal{S}
 - (b) **Rank** $(x_i, y_i) \in \mathcal{D}_d$ based on these scores
 - (c) **Find** $\mathcal{D}_s \subset \mathcal{D}_d$ by selecting top $p\%$ of \mathcal{D}_d
 - (d) **for** n_epochs **do**
| Fine-tune \mathcal{M} on \mathcal{D}_s
- end**
-

better subset \mathcal{D}_s by selecting top $p\%$ sentence pairs from the ranked \mathcal{D}_d . Finally, we fine-tune the base model \mathcal{M} on \mathcal{D}_s . LASER and DCCE performs *denoising* curriculum (*i.e.*, higher rank for good sentence pairs and lower rank for noisy ones) while MML performs *domain similarity* curriculum on the given data. Algorithm 1 presents a pseudo-code of the deterministic curriculum strategy.

2.2 Online Curriculum

Unlike deterministic curriculum, in this strategy the selected data subset \mathcal{D}_s changes dynamically in each epoch of the *model fine-tuning* stage through instantaneous feedback from the current NMT model. Specifically, in each epoch, we rank $(x_i, y_i) \in \mathcal{D}_d$ by leveraging the prediction scores from the emerging NMT model which assigns a probability to each token in the target sentence y_i . We then take the average of the token-level log probabilities to get the sentence-level probability score $P_\theta(y_i|x_i)$ which is regarded as the *prediction score* for the sentence pair (x_i, y_i) . Formally,

$$P_\theta(y_i|x_i) = \frac{1}{\ell} \sum_{t=1}^{\ell} \log p_\theta(y_{i,t}|y_{i,<t}, x_i) \quad (5)$$

This prediction score indicates the *confidence* of the emerging NMT model to generate the target sentence y_i from the source sentence x_i . Intuitively, if the model can predict the target sentence of a training data sample (x_i, y_i) with higher confidence, it indicates that the sample is *too easy* for the model and might not contain useful information to improve the NMT model further at that state. On the other hand, if a target sentence is predicted with lower confidence, it indicates that the training data

Algorithm 2 Online Curriculum Strategy

Input : General corpus \mathcal{D}_g , in-domain corpus $\mathcal{D}_d \subseteq \mathcal{D}_g$

Output : A trained translation model

1. // model warm-up stage
Train a base model \mathcal{M} on general domain corpus \mathcal{D}_g for K number of updates
 2. // model fine-tuning stage
for n_epochs **do**
 - (a) **Score** each $(x_i, y_i) \in \mathcal{D}_d$ by emerging \mathcal{M}
 - (b) **Rank** $(x_i, y_i) \in \mathcal{D}_d$ based on these scores
 - (c) **Find** $\mathcal{D}_s \subset \mathcal{D}_d$ by picking a *data-selection window*
 - (d) **Fine-tune** \mathcal{M} on \mathcal{D}_s
- end**
-

sample might be *too hard* for the model at that state, or it might be a noisy sample. Subsequently, including such hard or noisy samples in training at that state might degrade the NMT model performance.

Algorithm 2 presents the pseudo-code of our online data-selection curriculum strategy. After the *model warm-up* stage, we fine-tune \mathcal{M} for n_epochs on data subset \mathcal{D}_s which is selected in every epoch based on the emerging NMT models' confidence. Specifically, in the beginning of each epoch in the *model fine-tuning* stage, we find the *prediction score* $P_\theta(y_i|x_i)$ of each sample $(x_i, y_i) \in \mathcal{D}_d$. We then rank \mathcal{D}_d based on these scores and select $\mathcal{D}_s \subset \mathcal{D}_d$ by picking a *data-selection window* in the ranked data. Finally, we fine-tune \mathcal{M} on \mathcal{D}_s for that epoch. We present the conceptual demonstration of our online curriculum strategy in Figure 1. For picking the data-selection window in ranked \mathcal{D}_d , we investigate two methods:

- **Static Data-selection Window** Here in each epoch, we discard a *constant* amount (%) of easy and hard/noisy samples from \mathcal{D}_d based on the prediction scores and select the rests as \mathcal{D}_s . Even though in this method the *size* of the selected data subset (\mathcal{D}_s) is constant through out the model fine-tuning stage, unlike deterministic strategy the samples in \mathcal{D}_s *varies* from epoch-to-epoch due to the change in their prediction scores by the emerging NMT model \mathcal{M} .¹

- **Dynamic Data-selection Window** Unlike the static approach, here we *change* the data-selection window size in subsequent epochs. This can be

¹We present an illustrative example in Appendix B.




Figure 1: Conceptual demonstration of online curriculum. We rank the sentence pairs based on the prediction scores of the emerging NMT model \mathcal{M} and pick a *data-selection window* which discards easy and hard/noisy ones.

done in two ways:

- (i) *Expansion*: Begin with a smaller window ($|\mathcal{D}_s| \ll |\mathcal{D}_d|$) and gradually increase the window to a maximum size λ_{max} .
- (ii) *Shrink*: Begin with a larger window ($|\mathcal{D}_s| \sim |\mathcal{D}_d|$) and gradually decrease the window to a minimum size λ_{min} .

To change the data-selection window size, we use *linear scheduler*² which can be regarded as a function $\lambda(t)$ to map the current training epoch t to a scalar value. This value is regarded as the data-selection window size at epoch t . Formally,

$$\lambda_{exp}(t) = \begin{cases} \lambda_{init} + l_{inc} * t, & \text{if } \lambda_{exp}(t) < \lambda_{max} \\ \lambda_{max}, & \text{otherwise} \end{cases} \quad (6)$$

$$\lambda_{shr}(t) = \begin{cases} \lambda_{init} - l_{dec} * t, & \text{if } \lambda_{shr}(t) > \lambda_{min} \\ \lambda_{min}, & \text{otherwise} \end{cases}$$

where λ_{init} is the initial window size which is smaller for *expansion* and larger for *shrink*, and l_{inc}, l_{dec} are the hyperparameters of the respective schedulers.

3 Experimental Setup

Datasets We conduct experiments on six language pairs: three *high-resource* including English (En) to/from German (De), Hungarian (Hu), and Estonian (Et); and three *low-resource* including English (En) to/from Hausa (Ha), Tamil (Ta), and Malay (Ms). We use the dataset provided in WMT 2021³ — De and Ha are from *News shared task*,

²We experiment with other schedulers (Appendix C).

³<http://www.statmt.org/wmt21/>

| Pair | Train | | Validation | Test |
|-------|------------|-----------|------------|------|
| | All-data | In-domain | | |
| En-De | 89,893,260 | 2,152,577 | 1997 | 1418 |
| De-En | 89,893,260 | 2,152,577 | 2000 | 785 |
| En-Hu | 53,219,023 | 647,106 | 997 | 1012 |
| En-Et | 19,685,308 | 869,537 | 997 | 1012 |
| En-Ms | 1,694,311 | – | 997 | 1012 |
| En-Ta | 1,064,032 | – | 997 | 1012 |
| En-Ha | 685,780 | – | 1000 | 1000 |

Table 1: Dataset statistics after cleaning and deduplication.

while the remaining four pairs are from *Large-Scale Multilingual MT shared task*. For En \leftrightarrow De, we use newstest2019 as validation set and report test results on newstest2020. For En \leftrightarrow Ha, we randomly split the provided dev set into validation and test set. For the other language pairs, we use the official evaluation data (dev and devtest) as validation and test sets. Table 1 presents the dataset statistics after cleaning and deduplication. For high-resource language pairs, we consider formal texts parallel data corpora sources as in-domain ($\mathcal{D}_d \subset \mathcal{D}_g$), while for low-resource pairs, we do not differentiate between general-domain and in-domain corpus ($\mathcal{D}_d := \mathcal{D}_g$). Table 2 shows the in-domain corpora sources for high-resource language pairs.

Model Settings We use the Transformer (Vaswani et al., 2017) implementation in Fairseq (Ott et al., 2019); details of our model architecture settings are given in Appendix A. We use sentencepiece library⁴ to learn joint Byte-Pair-Encoding (BPE) of size 32,000 and 16,000 for En \leftrightarrow De and En \leftrightarrow Ha, respectively. For other language pairs, we use the official sentencepiece model provided in *Large-Scale Multilingual MT shared task*. We filter out parallel data with a length longer than 250 tokens during training. All experiments are evaluated using SacreBLEU (Post, 2018).

For LM training in the modified Moore-Lewis method (§2.1), we use the implementation in Fairseq. For in-domain LM training, we use 5M sentences from newscrawl, while we combine 10M commoncrawl data with newscrawl totaling 15M sentences to train the general-domain LM.

Baselines We compare our methods with the **converged model**, which is a standard NMT model trained on all the general-domain data (\mathcal{D}_g) until convergence. Additionally, we compare both

⁴<https://github.com/google/sentencepiece>

| Pair | In-domain Corpora |
|-------|--|
| En-De | Europarl, News Commentary |
| En-Hu | EUconst, Europarl, GlobalVoices, Wikipedia, WikiMatrix, WMT-News |
| En-Et | EUconst, Europarl, WikiMatrix, WMT-News |

Table 2: In-domain corpora sources for high-resource language pairs.

the deterministic and online curriculum approaches with the **traditional fine-tuning** approach where we fine-tune the base model from the warm-up stage with all the in-domain train data (\mathcal{D}_d) until convergence.

4 Results

The main results for the low- and high-resource languages are shown in Tables 3 and 4, respectively. For low-resource languages, we train the warm-up stage models for 20K updates, while the converged models are trained for approximately 50K updates. For high-resource languages, we train for 50K and 100K updates for the warm-up and converged models, respectively. In traditional fine-tuning (*Traditional Ft.* row in the Tables), we use all the available in-domain data (\mathcal{D}_d) in each fine-tuning epoch. On the other hand, for both deterministic and online curricula, we use *at most* 40% of the available in-domain data ($\mathcal{D}_s \subset \mathcal{D}_d$) in each fine-tuning epoch.

Comparing the performance of traditional fine-tuning with the *Converged Model* on low-resource languages (Table 3), we see that both of these perform on par. This is not surprising as both approaches use all the data (\mathcal{D}_g) during the whole training (for low-resource languages $\mathcal{D}_d := \mathcal{D}_g$). The only difference between the two approaches is – while the converged model continues to train the base model from the warm-up stage, the traditional fine-tuning approach resets the base model’s meta-parameters (*e.g.*, learning-rate, lr-scheduler, data-loader, optimizer) and continue the training.

For high-resource languages (Table 4), we fine-tune the base model only on the in-domain training data ($\mathcal{D}_d \subset \mathcal{D}_g$) in traditional fine-tuning, while the converged model continues to train the base model on all the general-domain data (\mathcal{D}_g). Here, traditional fine-tuning performs better than the converged model on En-De (+0.4) and En-Et (+0.9) but exhibits poor performance on the other four directions by 0.7 BLEU score on an average.

In the following, we discuss the performance of our data-selection curriculum approaches:

4.1 Performance of Deterministic Curricula

First, we consider the performance of deterministic curriculum approaches on low-resource languages. From Table 3, we see that fine-tuning the base model on the data subset (\mathcal{D}_s) selected by LASER outperforms the baseline (*Converged Model*) on five out of six translation tasks with a +2.2 BLEU gain in Ha-En. For the other two scoring methods, dual conditional cross-entropy (DCCE) and modified Moore-Lewis (MML), we also see a better or similar performance on 5/6 translation tasks. Compared to the traditional fine-tuning, the deterministic approaches perform better in most of the tasks – on average +0.5, +0.4, +0.2 BLEU gains for LASER, DCCE, and MML, respectively.

In Table 4, we see a similar trend of better performance of the deterministic curricula over the converged model on high-resource languages. Specifically, fine-tuning on the data subset selected by utilizing the scoring of both LASER and DCCE performs better on four out of six translation tasks, while the MML-based method achieves better performances on three tasks. The margins of improved performances for the high-resource languages are *higher* compared to the low-resource languages: +1.4, +0.9, +0.7 BLEU gains on average for DCCE, LASER, and MML, respectively over the baseline. If we compare with the traditional fine-tuning, the deterministic curriculum approaches perform better in most of the tasks – on average +1.2, +0.8, +0.4 BLEU better for DCCE, LASER, and MML, respectively.

To observe the better performance of the deterministic curriculum approaches more clearly, we fine-tune the base model from the warm-up stage with different percentages of ranked parallel data selected by the bitext scoring methods. Figure 2 shows the results. We observe that there exist multiple subsets of data ($\mathcal{D}_s \subset \mathcal{D}_d$), fine-tuning the base model on which demonstrates better performance compared to the *Converged Model* and *traditional fine-tuning*. For De-En, traditional fine-tuning (on 100% data) reduces BLEU score by 0.3 from the base model, while fine-tuning on most of the subsets selected by the deterministic curricula leads to improved performances. For Hu-En, traditional fine-tuning diminishes the performance of the base model by 0.5 BLEU. Unlike De-En, here we could

| Type | Setting | % data-used in each ep. | En-Ha | | En-Ms | | En-Ta | |
|--|------------------------|----------------------------|-----------------------------|-----------------------------|-----------------------------|-----------------------------|----------------------------|-----------------------------|
| | | | → | ← | → | ← | → | ← |
| Warm-up Model | All Data | 100% | 13.5 | 14.7 | 30.8 | 27.3 | 8.5 | 15.4 |
| Converged Model | All Data | 100% | 14.3 | 15.3 | 31.4 | 27.9 | 8.8 | 15.7 |
| <i>Warm-up Stage Model Fine-tuning (Ft.)</i> | | | | | | | | |
| Traditional Ft. | All Data | 100% | 14.4 ^{+0.1} | 15.6 ^{+0.3} | 31.5 ^{+0.1} | 28.0 ^{+0.1} | 8.7 ^{-0.1} | 15.7 ^{+0.0} |
| Det. Curricula | LASER | 40% | 14.6 ^{+0.3} | 17.5 ^{+2.2} | 31.7 ^{+0.3} | 28.2 ^{+0.3} | 8.8 ^{+0.0} | 15.9 ^{+0.2} |
| | Dual Cond. CE (DCCE) | 40% | 14.3 ^{+0.0} | 16.3 ^{+1.1} | 31.4 ^{+0.0} | 28.2 ^{+0.3} | 8.6 ^{-0.2} | 16.0 ^{+0.3} |
| | Mod. Moore-Lewis (MML) | 40% | 14.8 ^{+0.5} | 15.6 ^{+0.3} | 31.6 ^{+0.2} | 28.1 ^{+0.2} | 9.0 ^{+0.2} | 15.6 ^{-0.1} |
| Online Curricula | Static Window | 40% | 14.7 ^{+0.4} | 16.1 ^{+0.8} | 31.6 ^{+0.2} | 28.3 ^{+0.4} | 9.1 ^{+0.3} | 16.2 ^{+0.5} |
| | Dynamic Window | | | | | | | |
| | Expansion | <40% | 14.9 ^{+0.6} | 16.6 ^{+1.3} | 31.8 ^{+0.4} | 28.4 ^{+0.5} | 9.2 ^{+0.4} | 16.2 ^{+0.5} |
| | Shrink | <40% | 14.7 ^{+0.4} | 15.9 ^{+0.6} | 31.4 ^{+0.0} | 28.3 ^{+0.4} | 8.8 ^{+0.0} | 16.0 ^{+0.3} |
| Det. + Online | Hybrid | 15-20% | 14.7 ^{+0.4} | 16.4 ^{+1.1} | 31.5 ^{+0.1} | 28.2 ^{+0.3} | 9.1 ^{+0.2} | 15.9 ^{+0.2} |

Table 3: Main results for **low-resource** languages. Here, the data-percentage represents *general-domain* data (\mathcal{D}_g) and we do not differentiate between general-domain and in-domain corpus ($\mathcal{D}_d := \mathcal{D}_g$). Subscript values denote the BLEU score differences from the respective converged model.

| Type | Setting | % data-used in each ep. | En-De | | En-Hu | | En-Et | |
|--|------------------------|----------------------------|-----------------------------|-----------------------------|-----------------------------|----------------------|-----------------------------|----------------------|
| | | | → | ← | → | ← | → | ← |
| Warm-up Model | All Data | 100%+OOD | 34.9 | 40.8 | 33.9 | 36.0 | 35.7 | 37.1 |
| Converged Model | All Data | 100%+OOD | 36.1 | 41.2 | 35.9 | 36.7 | 36.7 | 38.2 |
| <i>Warm-up Stage Model Fine-tuning (Ft.)</i> | | | | | | | | |
| Traditional Ft. | All In-domain Data | 100% | 36.5 ^{+0.4} | 40.5 ^{-0.7} | 35.4 ^{-0.5} | 35.5 ^{-1.2} | 37.6 ^{+0.9} | 37.4 ^{-0.8} |
| Det. Curricula | LASER | 40% | 37.6 ^{+1.5} | 42.4 ^{+1.2} | 36.0 ^{+0.1} | 35.9 ^{-0.8} | 37.6 ^{+0.9} | 37.8 ^{-0.4} |
| | Dual Cond. CE (DCCE) | 40% | 37.9 ^{+1.8} | 43.0 ^{+1.8} | 36.2 ^{+0.3} | 35.4 ^{-1.3} | 38.0 ^{+1.3} | 37.3 ^{-0.9} |
| | Mod. Moore-Lewis (MML) | 40% | 37.1 ^{+1.0} | 41.7 ^{+0.5} | 35.8 ^{-0.1} | 35.2 ^{-1.5} | 37.3 ^{+0.6} | 37.4 ^{-0.8} |
| Online Curricula | Static Window | 40% | 37.3 ^{+1.2} | 41.4 ^{+0.2} | 36.1 ^{+0.2} | 35.4 ^{-1.3} | 37.9 ^{+1.2} | 37.7 ^{-0.5} |
| | Dynamic Window | | | | | | | |
| | Expansion | <40% | 37.3 ^{+1.2} | 41.6 ^{+0.4} | 36.4 ^{+0.5} | 35.6 ^{-1.1} | 38.1 ^{+1.4} | 37.8 ^{-0.4} |
| | Shrink | <40% | 37.0 ^{+0.9} | 41.2 ^{+0.0} | 36.0 ^{+0.1} | 35.7 ^{-1.0} | 38.0 ^{+1.3} | 37.6 ^{-0.6} |
| Det. + Online | Hybrid | 15-20% | 38.1 ^{+2.0} | 43.3 ^{+2.1} | 36.1 ^{+0.2} | 35.6 ^{-1.1} | 37.9 ^{+1.2} | 37.3 ^{-0.9} |

Table 4: Main results for **high-resource** languages. Here, the data-percentage represents only *In-domain* data (\mathcal{D}_d) from Table 1 and *100%+OOD* denotes *All-data* (\mathcal{D}_g). Subscript values denote the BLEU score differences from respective converged model.

not find a subset by the deterministic curricula fine-tuning on which improves the performance of the base model.

4.2 Performance of Online Curricula

Our online curriculum approaches perform on par with the deterministic curricula for both low- and high-resource languages as shown in Tables 3 and 4, respectively. Unlike deterministic approaches, here we leverage the emerging models’ prediction scores without using any external pretrained scoring methods. In our static window approach, we discard the top 30% and bottom 30% sentence pairs from the ranked \mathcal{D}_d and fine-tune the base model from the warm-up stage on the remaining 40% data (\mathcal{D}_s). The selected data in \mathcal{D}_s vary dynamically from epoch-to-epoch due to the change in the prediction scores of the emerging NMT models (Figure 4). From the results (Tables 3, 4), we notice that

the data-selection by *Static Window* method outperforms the *Converged Model* on ten out of twelve translation tasks and the BLEU scores are comparable to the deterministic curriculum approaches.

In our dynamic window approach, we either expand or shrink the window size, where the selected window is confined to the range of 30% to 70% of the ranked \mathcal{D}_d , *i.e.*, \mathcal{D}_s can be at most 40% of \mathcal{D}_d . In window *expansion*, we start \mathcal{D}_s with 10% of \mathcal{D}_d and linearly increase it to 40% in the subsequent epochs, while in the window *shrink* method we start \mathcal{D}_s with 40% and linearly decrease to 10% of \mathcal{D}_d . With dynamic window *expansion*, we achieve slightly better (up to +0.5 BLEU) performance on ten out of twelve translation tasks compared to the static window method. On the other hand, the dynamic window *shrink* method performs slightly lower than window expansion in most of the translation tasks.




Figure 2: Fine-tuned *warm-up stage model* on different sizes of ranked data (deterministic curricula).

5 Discussion and Analysis

5.1 Hybrid Curriculum

To benefit from both deterministic and online curricula, we combine the two strategies. Specifically, we consider three subsets of data comprising of the top 50% of \mathcal{D}_d ranked by each of the three bitext scoring methods in §2.1 and keep the common sentence pairs (intersection of three subsets). We then apply the static window data-selection curriculum on these sentence pairs, where we discard the top 10% and bottom 10% pairs (ranked by the emerging model’s prediction scores) and fine-tune the base model from the warm-up stage on the remaining bitext. Depending on the language pairs, the data percentage for the fine-tuning stage (\mathcal{D}_s) becomes 15-20% of \mathcal{D}_d . Despite being a smaller subset of data for fine-tuning, performances of the hybrid curriculum strategy are better on 10 out of 12 translation tasks compared to the baseline (Table 3, 4). Notably, for En-De and De-En, the hybrid curriculum attains +2.0 and +2.1 BLEU scores compared to the converged model.

5.2 Are All Data Useful Always?

Our proposed training framework uses all the data (\mathcal{D}_g) in the model warm-up stage and then utilizes subsets of in-domain data (\mathcal{D}_s) in the model fine-tuning stage. This resembles the “*formal education system*” where students first learn the general subjects with the same weights and later concen-

| Corpus | En-De | | En-Hu | | En-Et | |
|-----------|-------|------|-------|------|-------|------|
| | → | ← | → | ← | → | ← |
| All-data | 36.1 | 41.2 | 35.9 | 36.7 | 36.7 | 38.2 |
| In-domain | 32.6 | 33.5 | 25.5 | 23.6 | 30.6 | 30.3 |

Table 5: Results for high-resource languages on **all-data** (\mathcal{D}_g) vs. **in-domain data** (\mathcal{D}_d) when trained from a random state until convergence.

trate more on a selected subset of specialized subjects. The first stage teaches them the base knowledge which is useful in the ensuing stage. We observe a similar phenomenon in our experiments. From Table 5, we see that the performance of the NMT model using only the in-domain data is worse than using all general-domain data (-8.1 BLEU on average). Moreover, our data-selection curriculum training framework outperforms the converged model that uses all the data throughout the training in most of the translation tasks by a sizable margin. This indicates that *not all data are useful all the time*. Additionally, Figure 2 shows that in most scenarios, fine-tuning on selected data subsets \mathcal{D}_s outperform the traditional fine-tuning that uses all the data. This observation validates our intuition that some data samples are not only redundant but also detrimental to the NMT model’s performance.

5.3 Do We Need the Two Stages?

For the online curricula, we leverage the model \mathcal{M} for selecting \mathcal{D}_s based on the prediction scores, while in the deterministic curricula, we do not use

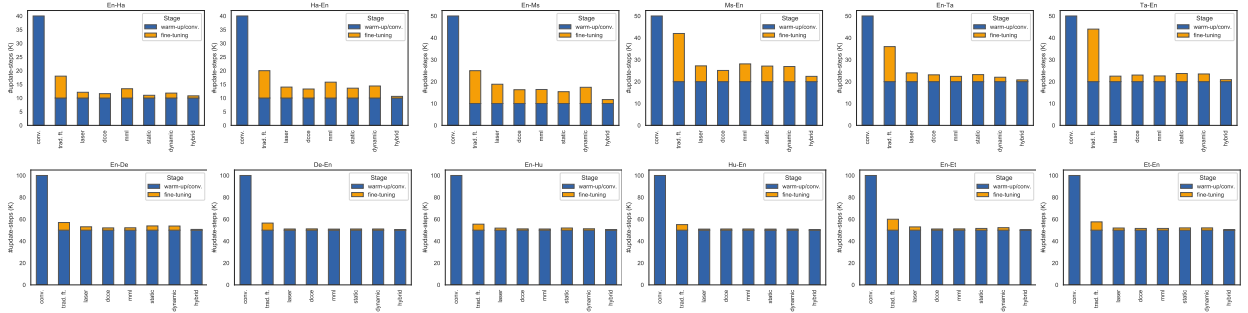


Figure 3: Number of update steps required for each setting of Tables 3, 4. We keep batch size same in each setting.

| Scoring Method | Top data% | En-Ha | | En-Ms | | En-Ta | |
|----------------------|-----------|----------------------|----------------------|----------------------|----------------------|--------------------|----------------------|
| | | → | ← | → | ← | → | ← |
| LASER | 10% | 14.1 _{8.3} | 17.3 _{10.1} | 30.9 _{18.9} | 27.9 _{15.1} | 8.1 _{0.7} | 15.8 _{1.6} |
| | 40% | 14.6 _{13.1} | 17.5 _{16.5} | 31.7 _{30.2} | 28.2 _{25.2} | 8.8 _{5.9} | 15.9 _{10.7} |
| Dual Cond. CE | 10% | 13.0 _{1.3} | 16.3 _{8.0} | 31.0 _{18.4} | 28.0 _{15.5} | 8.0 _{0.0} | 15.2 _{0.2} |
| | 40% | 14.3 _{12.9} | 16.3 _{15.3} | 31.4 _{29.5} | 28.2 _{25.0} | 8.6 _{5.3} | 16.0 _{11.0} |
| Modified Moore-Lewis | 10% | 14.4 _{5.9} | 15.1 _{4.7} | 31.8 _{19.6} | 27.9 _{15.3} | 8.5 _{0.0} | 15.2 _{0.6} |
| | 40% | 14.8 _{13.3} | 15.6 _{13.6} | 31.6 _{30.8} | 28.1 _{24.9} | 9.0 _{5.9} | 15.6 _{10.5} |

Table 6: Results for **two-stage training framework vs. training without warm-up stage** on top 10% and 40% of selected data ranked by three scoring methods (§2.1). Main values denote the results of our two-stage framework utilizing warm-up stage, while subscript values represent results when model is trained on the same data subset without warm-up stage.

the emerging model for selecting the data subset. One might ask – *do we need a base model in the deterministic curricula? Can we get rid of the warm-up stage?* To answer these questions, we perform another set of experiments where we train \mathcal{M} from a randomly initialized state on the top $p\%$ of the selected data ($p = \{10, 40\}$) ranked by the three bitext scoring methods (§2.1) and compare the results with our two-stage training framework where we fine-tune the base model from the warm-up stage on the same data subset. From the results in Table 6, it is evident that our proposed training framework utilizing the warm-up stage outperforms the approach not using any warm-up stage by a sizable margin in all the tasks.

5.4 Comparing Required Update Steps

Our proposed curriculum training approaches not only exhibit better performance but also converge faster compared to the baseline and traditional fine-tuning method. In Figure 3, we plot the number of update steps required by each of the settings in Table 3 and 4. On average, we need about 50% fewer updates compared to the converged model. For high-resource languages, we need much fewer updates in the model fine-tuning stage. For all the

| Type | Setting | %data-used in each ep. | En-De | |
|--|------------------------|------------------------|--------------------|--------------------|
| | | | → | ← |
| Warm-up Model | All Data | 100% | 33.3 | 39.1 |
| Converged Model | All Data | 100% | 34.6 | 40.0 |
| <i>Warm-Up Model Fine-tuning (Fr.)</i> | | | | |
| Traditional Ft. | All data | 100% | 34.0 -0.6 | 41.6 $+1.6$ |
| Det. Curricula | LASER | 40% | 34.4 -0.2 | 43.2 $+3.2$ |
| | Dual Cond. CE (DCCE) | 40% | 35.1 $+0.5$ | 44.4 $+4.4$ |
| | Mod. Moore-Lewis (MML) | 40% | 34.5 -0.1 | 41.6 $+1.6$ |
| Online Curricula | Static Window | 40% | 34.1 -0.5 | 41.9 $+1.9$ |
| | Dynamic Window | | | |
| | Expansion | <40% | 34.4 -0.2 | 42.2 $+2.2$ |
| Shrink | <40% | 34.3 -0.3 | 42.0 $+2.0$ | |

Table 7: Results for En \leftrightarrow De on **noisy ParaCrawl corpus** of 10M bitext pairs. Here, the data-percentage corresponds to all 10M bitext (\mathcal{D}_g) and $\mathcal{D}_d := \mathcal{D}_g$. Subscript values denote the BLEU score difference from the respective converged model.

language pairs, the hybrid curriculum strategy requires the fewest updates as the size of selected subsets is much lower compared to other approaches.

5.5 Performance on Noisy Data

We further evaluate our framework on noisy data. We randomly selected 10M bitext pairs from the En-De ParaCrawl corpus (Bañón et al., 2020). We keep the experimental settings similar to §4 and present the results in Table 7. Fine-tuning on the data subset (\mathcal{D}_s) selected by DCCE method outperforms the baseline (*Converged Model*) on both directions with a +4.4 BLEU gain in De-En. All the other deterministic and online curriculum methods perform better than the converged model on the De-En direction with a sizable margin. Compared to the traditional fine-tuning, all the curriculum methods perform better in both En to/from De.

6 Related Work

Curriculum Learning Inspired by human learners, Elman (1993) argues that optimization of neural network training can be accelerated by gradually increasing the difficulty of the concepts. Bengio et al. (2009) were the first to use the term “curricu-

lum learning” to refer to the easy-to-hard training strategies in the context of machine learning. Using an easy-to-hard curriculum based on increasing vocabulary size in language model training, they achieved performance improvement. Recent work (Jiang et al., 2015; Hach Cohen and Weinshall, 2019; Zhou et al., 2020a) shows that manoeuvring the sequence of training data can improve both training efficiency and model accuracy. Several studies show the effectiveness of the difficulty-based curriculum learning in a wide range of NLP tasks including task-specific word representation learning (Tsvetkov et al., 2016), natural language understanding tasks (Sachan and Xing, 2016; Xu et al., 2020a), reading comprehension (Tay et al., 2019), and language modeling (Campos, 2021). Several studies show the effectiveness of the difficulty-based curriculum learning in a wide range of NLP tasks (Cirik et al., 2016; Liu et al., 2018).

Curriculum Learning in NMT The difficulty-based curriculum in NMT was first explored by Kocmi and Bojar (2017). Later, Zhang et al. (2018) adopt a probabilistic view of curriculum learning and investigate a variety of difficulty criteria based on human intuition, *e.g.*, sentence length and word rarity. Platanios et al. (2019) connect the appearance of difficult samples with NMT model competence. Liu et al. (2020) propose a norm-based curriculum learning method based on the norm of word embedding. Zhou et al. (2020b) use a pre-trained language model to measure the sentence-level uncertainty. Zhan et al. (2021) propose meta-curriculum learning for domain adaptation in NMT. Zhou et al. (2021) adopt sentence-level BLEU score as difficulty measure. Most curriculum learning methods in NMT focus on addressing the batch selection issue from the beginning of the training by using hand-designed heuristics (Zhao et al., 2020). In contrast, our proposed two-stage training framework for NMT fine-tunes the base model from the warm-up stage on *selected* subsets of data. Our data-selection curriculum training framework is more realistic, resembling the formal education system as discussed in §5.2.

Self-paced Learning in NMT Here, the model itself measures the difficulty of the training samples to adjust the learning pace (Kumar et al., 2010). In their approach, Wan et al. (2020) first train the NMT model for M passes on the data and cache the translation probabilities to find variance. The lower variance of the translation probabilities of a

sample reflects higher confidence. Later, they use the confidence scores as factors to weight the loss to control the model updates. For low-resource NMT, Xu et al. (2020b) utilize the declination of the loss of a sample as the difficulty measure and train the model on easier samples (higher loss drop). In our online curriculum, we leverage the prediction scores of the emerging model in the model fine-tuning stage. However, after ranking the samples based on the prediction scores, we employ a *variety* of data-selection methods to select the better data subset (§2.2).

Domain Specific Fine-tuning in NMT Here, converged NMT models trained on large general-domain parallel data are fine-tuned on all in-domain data (Luong and Manning, 2015; Zoph et al., 2016; Freitag and Al-Onaizan, 2016). In contrast, in our framework, we adapt a base NMT model (non-converged) on *selected* subsets of the in-domain data considering the data usefulness and quality.

Data-selection Strategy in NMT We apply the data-selection curriculum in the model fine-tuning stage in our framework. Joty et al. (2015) use domain adaptation by penalizing sequences similar to the out-domain data in their training data-selection. van der Wees et al. (2017) gradually decrease the training data size to a cleaner subset of the data estimated by some external scorers. Wang et al. (2018) propose a curriculum-based data-selection strategy by using an *additional* trusted clean dataset to calculate the noise level of a sample. Kumar et al. (2019) use reinforcement learning to learn a denoising curriculum jointly with the NMT system. Jiao et al. (2020) identify the inactive samples during training and re-label them for later use. Wang et al. (2021) find gradient alignments between a clean dataset and the training data to mask out noisy data.

7 Conclusion

We have presented a two-stage training framework for NMT where we apply a data-selection curriculum in the model fine-tuning stage. Our novel online curriculum strategy utilizes the emerging models’ prediction scores for the selection of a better data subset. Experiments on six low- and high-resource language pairs show the efficacy of our proposed framework. Our curriculum training approaches exhibit better performance as well as converge much faster by requiring fewer updates compared to the baselines.

Limitations

Despite its effectiveness, our proposed data selection curriculum has some limitations:

- The deterministic curriculum (§2.1) uses external scoring methods that require pretrained models for selecting the data subset to be used in the model fine-tuning stage. These scoring methods' pretraining incur additional training costs.
- Our online curriculum approach (§2.2) on the other hand, is free from such additional pretraining costs. Nevertheless, it requires an extra forward propagation step in each of the epochs of the model fine-tuning stage to find the prediction scores of each of the sentence pairs. One possible way to avoid the extra forward propagation step is to cache the prediction scores while calculating the training losses in the previous epoch. However, there will be discrepancies in the prediction scores of samples as the predicted model will not be the same; it will get updated. Even though we did not investigate this avenue in our paper, we believe this can be an interesting research direction to observe the phenomenon.

References

- Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydrin, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(WMT21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88, Online. Association for Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019. [Margin-based parallel corpus mining with multilingual sentence embeddings](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3197–3203, Florence, Italy. Association for Computational Linguistics.
- Amittai Axelrod, Xiaodong He, and Jianfeng Gao. 2011. [Domain adaptation via pseudo in-domain data selection](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 355–362, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel L. Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, Sergio Ortiz Rojas, Leopoldo Pla Sempere, Gema Ramírez-Sánchez, Elsa Sarrías, Marek Strelec, Brian Thompson, William Waites, Dion Wiggins, and Jaime Zaragoza. 2020. [ParaCrawl: Web-scale acquisition of parallel corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567, Online. Association for Computational Linguistics.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 41–48, New York, NY, USA. Association for Computing Machinery.
- Daniel Campos. 2021. [Curriculum learning for language modeling](#). *CoRR*, abs/2108.02170.
- Volkan Cirik, Eduard H. Hovy, and Louis-Philippe Morency. 2016. [Visualizing and understanding curriculum learning for long short-term memory networks](#). *CoRR*, abs/1611.06204.
- Alexis Conneau, Guillaume Lample, Marc'Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. [Word translation without parallel data](#). *arXiv preprint arXiv:1710.04087*.
- Jeffrey L. Elman. 1993. [Learning and development in neural networks: the importance of starting small](#). *Cognition*, 48(1):71–99.
- Markus Freitag and Yaser Al-Onaizan. 2016. [Fast domain adaptation for neural machine translation](#). *CoRR*, abs/1612.06897.
- Guy Hach Cohen and Daphna Weinshall. 2019. [On the power of curriculum learning in training deep networks](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2535–2544. PMLR.
- Hany Hassan, Anthony Aue, C. Chen, Vishal Chowdhary, J. Clark, C. Federmann, Xuedong Huang, Marcin Junczys-Dowmunt, W. Lewis, M. Li, Shujie Liu, T. Liu, Renqian Luo, Arul Menezes, Tao Qin, F. Seide, Xu Tan, Fei Tian, Lijun Wu, Shuangzhi Wu, Yingce Xia, Dongdong Zhang, Zhirui Zhang, and M. Zhou. 2018. [Achieving human parity on automatic chinese to english news translation](#). *ArXiv*, abs/1803.05567.
- Lu Jiang, Deyu Meng, Qian Zhao, Shiguang Shan, and Alexander G. Hauptmann. 2015. [Self-paced curriculum learning](#). In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15*, page 2694–2700. AAAI Press.

- Wenxiang Jiao, Xing Wang, Shilin He, Irwin King, Michael Lyu, and Zhaopeng Tu. 2020. [Data Rejuvenation: Exploiting Inactive Training Examples for Neural Machine Translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2255–2266, Online. Association for Computational Linguistics.
- Shafiq Joty, Hassan Sajjad, Nadir Durrani, Kamla Al-Mannai, Ahmed Abdelali, and Stephan Vogel. 2015. [How to avoid unwanted pregnancies: Domain adaptation using neural network models](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1259–1270, Lisbon, Portugal. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt. 2018. [Dual conditional cross-entropy filtering of noisy parallel corpora](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 888–895, Belgium, Brussels. Association for Computational Linguistics.
- Huda Khayrallah and Philipp Koehn. 2018. [On the impact of various types of noise on neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 74–83, Melbourne, Australia. Association for Computational Linguistics.
- Tom Kocmi and Ondřej Bojar. 2017. [Curriculum learning and minibatch bucketing in neural machine translation](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pages 379–386, Varna, Bulgaria. INCOMA Ltd.
- Gaurav Kumar, George Foster, Colin Cherry, and Maxim Krikun. 2019. [Reinforcement learning based curriculum optimization for neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2054–2061, Minneapolis, Minnesota. Association for Computational Linguistics.
- M. Kumar, Benjamin Packer, and Daphne Koller. 2010. [Self-paced learning for latent variable models](#). In *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc.
- Cao Liu, Shizhu He, Kang Liu, and Jun Zhao. 2018. [Curriculum learning for natural answer generation](#). In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 4223–4229. International Joint Conferences on Artificial Intelligence Organization.
- Xuebo Liu, Houtim Lai, Derek F. Wong, and Lidia S. Chao. 2020. [Norm-based curriculum learning for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 427–436, Online. Association for Computational Linguistics.
- Minh-Thang Luong and Christopher Manning. 2015. [Stanford neural machine translation systems for spoken language domains](#). In *Proceedings of the 12th International Workshop on Spoken Language Translation: Evaluation Campaign*, pages 76–79, Da Nang, Vietnam.
- Robert C. Moore and William Lewis. 2010. [Intelligent selection of language model training data](#). In *Proceedings of the ACL 2010 Conference Short Papers*, pages 220–224, Uppsala, Sweden. Association for Computational Linguistics.
- Elissa L. Newport. 1990. [Maturational constraints on language learning](#). *Cognitive Science*, 14(1):11–28.
- 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 *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Emmanouil Antonios Platanios, Otilia Stretcu, Graham Neubig, Barnabas Poczos, and Tom Mitchell. 2019. [Competence-based curriculum learning for neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1162–1172, Minneapolis, Minnesota. Association for Computational Linguistics.
- Martin Popel, Marketa Tomkova, Jakub Tomek, Łukasz Kaiser, Jakob Uszkoreit, Ondřej Bojar, and Zdeněk Žabokrtský. 2020. [Transforming machine translation: a deep learning system reaches news translation quality comparable to human professionals](#). *Nature Communications*, 11(1):1–15.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Mrinmaya Sachan and Eric Xing. 2016. [Easy questions first? a case study on curriculum learning for question answering](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 453–463, Berlin, Germany. Association for Computational Linguistics.
- Yi Tay, Shuohang Wang, Anh Tuan Luu, Jie Fu, Minh C. Phan, Xingdi Yuan, Jinfeng Rao, Siu Cheung Hui, and Aston Zhang. 2019. [Simple and effective curriculum pointer-generator networks for reading comprehension over long narratives](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4922–4931, Florence, Italy. Association for Computational Linguistics.

- Yulia Tsvetkov, Manaal Faruqui, Wang Ling, Brian MacWhinney, and Chris Dyer. 2016. [Learning the curriculum with Bayesian optimization for task-specific word representation learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 130–139, Berlin, Germany. Association for Computational Linguistics.
- Marlies van der Wees, Arianna Bisazza, and Christof Monz. 2017. [Dynamic data selection for neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1400–1410, Copenhagen, Denmark. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30, pages 5998–6008. Curran Associates, Inc.
- Yu Wan, Baosong Yang, Derek F. Wong, Yikai Zhou, Lidia S. Chao, Haibo Zhang, and Boxing Chen. 2020. [Self-paced learning for neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1074–1080, Online. Association for Computational Linguistics.
- Wei Wang, Taro Watanabe, Macduff Hughes, Tetsuji Nakagawa, and Ciprian Chelba. 2018. [Denoising neural machine translation training with trusted data and online data selection](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 133–143, Brussels, Belgium. Association for Computational Linguistics.
- Xinyi Wang, Ankur Bapna, Melvin Johnson, and Orhan Firat. 2021. [Gradient-guided loss masking for neural machine translation](#). *CoRR*, abs/2102.13549.
- Benfeng Xu, Licheng Zhang, Zhendong Mao, Quan Wang, Hongtao Xie, and Yongdong Zhang. 2020a. [Curriculum learning for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6095–6104, Online. Association for Computational Linguistics.
- Chen Xu, Bojie Hu, Yufan Jiang, Kai Feng, Zeyang Wang, Shen Huang, Qi Ju, Tong Xiao, and Jingbo Zhu. 2020b. [Dynamic curriculum learning for low-resource neural machine translation](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 3977–3989, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Runzhe Zhan, Xuebo Liu, Derek F. Wong, and Lidia S. Chao. 2021. [Meta-curriculum learning for domain adaptation in neural machine translation](#). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 14310–14318.
- Xuan Zhang, Gaurav Kumar, Huda Khayrallah, Kenton Murray, Jeremy Gwinnup, Marianna J Martindale, Paul McNamee, Kevin Duh, and Marine Carpuat. 2018. [An empirical exploration of curriculum learning for neural machine translation](#). *arXiv preprint arXiv:1811.00739*.
- Mingjun Zhao, Haijiang Wu, Di Niu, and Xiaoli Wang. 2020. [Reinforced curriculum learning on pre-trained neural machine translation models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):9652–9659.
- Lei Zhou, Liang Ding, Kevin Duh, Shinji Watanabe, Ryohhei Sasano, and Koichi Takeda. 2021. [Self-guided curriculum learning for neural machine translation](#). In *Proceedings of the 18th International Conference on Spoken Language Translation (IWSLT 2021)*, pages 206–214, Bangkok, Thailand (online). Association for Computational Linguistics.
- Tianyi Zhou, Shengjie Wang, and Jeffrey Bilmes. 2020a. [Curriculum learning by dynamic instance hardness](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 8602–8613. Curran Associates, Inc.
- Yikai Zhou, Baosong Yang, Derek F. Wong, Yu Wan, and Lidia S. Chao. 2020b. [Uncertainty-aware curriculum learning for neural machine translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6934–6944, Online. Association for Computational Linguistics.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575, Austin, Texas. Association for Computational Linguistics.

Appendix

A Model Architecture Settings

For En↔Ha, we use a smaller Transformer architecture with five layers, while for the other language pairs we use larger Transformer architecture with six encoder and decoder layers. We present the number of attention heads, embedding dimension, and the inner-layer dimension of both settings in Table 8.

| Settings | En↔Ha | Other Pairs |
|-----------------------|-------|-------------|
| Transformer Layers | 5 | 6 |
| #Attention Heads | 8 | 16 |
| Embedding Dimension | 512 | 1024 |
| Inner-layer Dimension | 2048 | 4096 |

Table 8: Model architecture settings.

B Variety of Data Samples in Static Data-selection Window

In the beginning of each epoch in the *model fine-tuning* stage in static data-selection window approach (§2.2), we rank \mathcal{D}_d based on the prediction scores of each sentence pair (x_i, y_i) . We then pick a *fixed* data-selection window (confined to a range of data-percentage in the ranked \mathcal{D}_d e.g., 30% to 70%) by discarding too easy and too hard/noisy samples. Even though in this approach the size of the selected data subset (\mathcal{D}_s) remains the same throughout the model fine-tuning stage, the samples in \mathcal{D}_s *changes* from epoch-to-epoch due to the change in their prediction scores by the current model. We present an illustrative example of this phenomenon in Figure 4.

In the *current* epoch of the fine-tuning stage (Figure 4(a)), samples 2, 3, 4, and 5 are selected to train the model while samples 1 and 6 are discarded – 1 is too hard/noisy and 6 is too easy for the current model. In the *next* epoch (Figure 4(b)), some samples might be selected again (samples 3 and 5), while some earlier selected samples might have lower prediction scores and not be selected due to the hardness to the current model (sample 2). Again, some previously selected samples might have higher prediction scores and not be selected due to the easiness (sample 4). And some samples not selected in the previous epoch can now be selected (samples 1 and 6).

| Scheduler | En-Ha | | En-De | |
|-------------|-------------|-------------|-------------|-------------|
| | → | ← | → | ← |
| Linear | 14.9 | 16.6 | 37.3 | 41.6 |
| Exponential | 14.5 | 16.3 | 36.7 | 41.0 |
| Square-Root | 14.4 | 16.0 | 36.9 | 40.9 |

Table 9: Results for En↔{Ha, De} for online curriculum with dynamic data-selection window (expansion) using different schedulers. We keep λ_{init} the same.

C Schedulers in Dynamic Data-selection Window

To change the data-selection window size in dynamic approach, we use schedulers which controls how the size of the window will grow in subsequent epochs (§2.2). Apart from the linear scheduler (Eq. 6), we also experiment with two other schedulers:

- **Exponential Scheduler** We find the data-selection window size (for window expansion and shrink) at epoch t using the following formula:

$$\lambda_{\text{exp}}^{\text{expn}}(t) = \begin{cases} \lambda_{\text{init}} * E_{\text{inc}}^t, & \text{if } \lambda_{\text{exp}}^{\text{expn}}(t) < \lambda_{\text{max}} \\ \lambda_{\text{max}}, & \text{otherwise} \end{cases} \quad (7)$$

$$\lambda_{\text{shr}}^{\text{expn}}(t) = \begin{cases} \lambda_{\text{init}} * E_{\text{dec}}^{-t}, & \text{if } \lambda_{\text{shr}}^{\text{expn}}(t) > \lambda_{\text{min}} \\ \lambda_{\text{min}}, & \text{otherwise} \end{cases}$$

where λ_{init} is the initial window size which is smaller for *expansion* and larger for *shrink*, and $E_{\text{inc}}, E_{\text{dec}}$ are the hyperparameters of the exponential schedulers.

- **Square-Root Scheduler** We find the data-selection window size (for window expansion and shrink) at epoch t using the following formula:

$$\lambda_{\text{exp}}^{\text{sqr}}(t) = \begin{cases} \sqrt{\lambda_{\text{init}}^2 + (C_1 - \lambda_{\text{init}}^2) * \frac{t}{S_{\text{inc}}}}, & \text{if } \lambda_{\text{exp}}^{\text{sqr}}(t) < \lambda_{\text{max}} \\ \lambda_{\text{max}}, & \text{otherwise} \end{cases}$$

$$\lambda_{\text{shr}}^{\text{sqr}}(t) = \begin{cases} \sqrt{\lambda_{\text{init}}^2 + (C_2 - \lambda_{\text{init}}^2) * \frac{t}{S_{\text{dec}}}}, & \text{if } \lambda_{\text{shr}}^{\text{sqr}}(t) > \lambda_{\text{min}} \\ \lambda_{\text{min}}, & \text{otherwise} \end{cases} \quad (8)$$

where λ_{init} is the initial window size which is smaller for *expansion* and larger for *shrink*, and $C_1, C_2, S_{\text{inc}}, S_{\text{dec}}$ are the hyperparameters of the square-root schedulers.

In our initial experiments, we explore the three schedulers — linear, exponential, and square-root.




Figure 4: Illustrative example of how data samples varies in *static data-selection window* approach in online curriculum. Even though the size of data-selection window is fixed throughout the model fine-tuning stage, the samples in the selected subsets vary from epoch-to-epoch due to the change in their prediction scores by the emerging model.

We found that linear scheduler performs better compared to the other schedulers. We present the results in Table 9.