

AUGVIC: Exploiting BiText Vicinity for Low-Resource NMT

Tasnim Mohiuddin^{*¶}, M Saiful Bari^{*¶}, and Shafiq Joty^{¶†}

[¶]Nanyang Technological University, Singapore

[†]Salesforce Research Asia, Singapore

{mohi0004, bari0001, srjoty}@ntu.edu.sg

Abstract

The success of Neural Machine Translation (NMT) largely depends on the availability of large bitext training corpora. Due to the lack of such large corpora in low-resource language pairs, NMT systems often exhibit poor performance. Extra relevant monolingual data often helps, but acquiring it could be quite expensive, especially for low-resource languages. Moreover, domain mismatch between bitext (train/test) and monolingual data might degrade the performance. To alleviate such issues, we propose AUGVIC, a novel data augmentation framework for low-resource NMT which exploits the vicinal samples of the given bitext without using any extra monolingual data explicitly. It can diversify the in-domain bitext data with finer level control. Through extensive experiments on four low-resource language pairs comprising data from different domains, we have shown that our method is comparable to the traditional back-translation that uses extra in-domain monolingual data. When we combine the synthetic parallel data generated from AUGVIC with the ones from the extra monolingual data, we achieve further improvements. We show that AUGVIC helps to attenuate the discrepancies between relevant and distant-domain monolingual data in traditional back-translation. To understand the contributions of different components of AUGVIC, we perform an in-depth framework analysis.

1 Introduction

Neural Machine Translation (NMT) has shown impressive performance in high-resource settings, even claiming to achieve parity with human professional translators (Hassan et al., 2018; Popel et al., 2020). Most successful NMT systems have billions of parameters (Lepikhin et al., 2021). They generally work well only when a good amount of parallel training data is available and perform poorly

in low-resource conditions (Koehn and Knowles, 2017; Guzmán et al., 2019). However, majority of the languages are low-resourced despite being used by large portion of world population. Hence, improving low-resource MT quality has been of great interests to the MT researchers.

There have been several attempts to extend the success of NMT in high-resource settings to low-resource language pairs that have a relatively small amount of available parallel data. Most of these methods mainly focus on leveraging extra monolingual data through back-translation (Sennrich et al., 2016) and self-training (He et al., 2020), or translation knowledge transfer through parallel data involving other assisting language pairs (Firat et al., 2016a,b; Johnson et al., 2017; Neubig and Hu, 2018).¹ Large scale pre-training is another recent trend to utilize large monolingual data for NMT (Liu et al., 2020). However, very few work has considered low-resource NMT without using auxiliary data or other pivot languages.

In the presence of a sufficient amount of in-domain monolingual data, back-translation (BT) has proved to be quite successful (Edunov et al., 2018). In this approach, a reverse intermediate model is trained on the original parallel data, which is later used to generate synthetic parallel data by translating sentences from target-side monolingual data into the source language. However, when there are scarcity of in-domain data which indeed a common situation in many low-resource settings, the success of BT may be limited (Chen et al., 2019).

Another understudied problem with BT is the issue with *domain mismatch* (Edunov et al., 2020). To elaborate, let us consider two scenarios: (i) the training and testing data come from the same or relevant domains (*e.g.*, News), and (ii) the test domain (News) is different from the training domain (*e.g.*,

^{*}Equal contribution

¹See (Dabre et al., 2020) for a survey of the later.

Subtitles). In the former case, we can foresee two problems. First, if we use out-of-domain monolingual data which is abundant, it might misguide the model and move it far away from the actual test distribution. Second, even if the monolingual data is from a domain similar to that of the training/testing data, there might be differences in topics, modality, style, etc., which might induce noise.

For the latter scenario, even if the monolingual data comes from the similar domain as the test data (News), the corresponding (reverse) translations will be noisy as the intermediate model would be trained on a different domain (Subtitles). Consequently, these noisy pseudo-parallel data will induce noise during training and might cause the model to perform worse (Wang et al., 2018). On the other hand, using in-domain (Subtitles) monolingual data in back-translation will not give enough diversity to cover the test domain (News).

In this work, inspired by the Vicinal Risk Minimization principle (Chapelle et al., 2001), we propose AUGVIC, a novel method to **augment vicinal** samples around the bitext distribution. Instead of using extra monolingual data, AUGVIC aims to leverage the vicinal samples of the original bitext, thereby enlarging the support of the training bitext distribution to improve model generalization. The main advantage is that the resulting distribution remains close to the original distribution and can be controlled at a finer level (Figure 1).

With the goal of training a source-to-target NMT system, AUGVIC augments vicinal samples in the target language. The vicinal samples are generated by predicting the masked tokens of a target bitext sentence using a pretrained large-scale language model. To generate synthetic bitext data from these augmented vicinal samples through a reverse intermediate (target-to-source) model, we propose two different methods: the first one is based on the traditional BT, while the second one leverages the original source sentence as a guide. Finally, we train the source-to-target model by combining the original parallel data with the synthetic bitext.

In order to demonstrate the effectiveness and robustness of AUGVIC, we conduct extensive experiments on four low-resource language pairs comprising data from different domains. Our results show significant improvements over the bitext baselines with 2.76 BLEU gains on an average on eight different translation tasks without using any extra monolingual data. AUGVIC also com-

plements traditional BT with additive gains when extra monolingual data is used. We also show AUGVIC’s efficacy in bridging the gap between in-domain and out-of-domain performance in traditional back-translation with monolingual data. We carried out an ablation study to understand the contribution of the diversity factor in our proposed framework. We open-source our framework at <https://ntunlp.sg.github.io/project/augvic/>.

2 Related Work

Two lines of studies are relevant to our work.

Low-resource NMT Although the main focus of investigation and improvement in NMT has been in high-resource settings, there has been a recent surge of interest in low-resource MT. However, achieving satisfactory performance in low-resource settings turns out to be challenging for NMT systems (Koehn and Knowles, 2017). Recent research has mainly focused on creating and cleaning parallel (Ramasamy et al., 2014; Islam, 2018) and comparable data (Tiedemann, 2012), utilizing bilingual lexicon induction (Conneau et al., 2017; Artetxe et al., 2018; Mohiuddin and Joty, 2019, 2020; Mohiuddin et al., 2020), fine-grained hyperparameter tuning (Sennrich and Zhang, 2019), and using other language pairs as pivot (Cheng et al., 2017; Kim et al., 2019).

Another avenue of research follows multilingual translation, where translation knowledge from high-resource language pairs are exploited by training a single NMT system on a mix of high-resource and low-resource language pairs (Firat et al., 2016a,b; Kocmi and Bojar, 2018; Gu et al., 2018; Neubig and Hu, 2018; Guzmán et al., 2019). Zoph et al. (2016) proposed a variant where they pretrain NMT system on a high-resource language pair before finetuning on a target low-resource language pair.

Data Augmentation for NMT Till now, one of the most successful data augmentation strategies in NMT is back-translation (BT) (Sennrich et al., 2016; Hoang et al., 2018), which exploits target-side monolingual data. Edunov et al. (2018) investigated BT extensively and scaled the method to millions of target-side monolingual sentences. Caswell et al. (2019) explored the role of noise in noised-BT and proposed to use a tag for back-translated source sentences. Besides BT, self-training is another data augmentation strategy for NMT which leverages source-side monolingual data (He et al.,

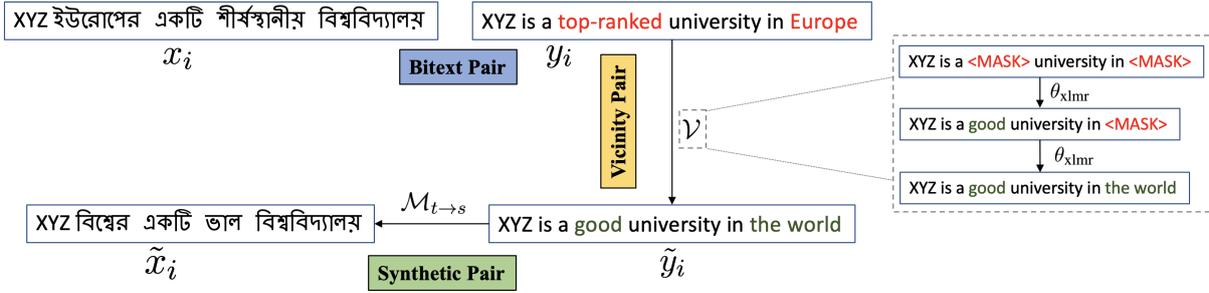


Figure 1: Illustration of AUGVIC steps for Bengali-to-English translation system. Here (x_i, y_i) is the original bitext pair, \tilde{y}_i is a vicinal sample of y_i , and $(\tilde{x}_i, \tilde{y}_i)$ is a synthetic pair where \tilde{x}_i is generated by a reverse intermediate translation system $\mathcal{M}_{t \rightarrow s}$. Right side of the figure shows the successive steps of vicinal sample generation.

2020). Large scale multilingual pre-training followed by bitext fine-tuning is a recent trend to utilize monolingual data for NMT, which is shown to be beneficial (Arivazhagan et al., 2019; Liu et al., 2020; Zhu et al., 2020; Lepikhin et al., 2021).

Apart from using extra monolingual data, Xie et al. (2017) show that data noising is an effective regularization method for NMT, while Wu et al. (2019) use noised training. In low-resource settings, Fadaee et al. (2017) augment bitext by replacing a common word with a low-frequency word in the target sentence, and change its corresponding word in the source sentence to improve the translation quality of rare words. Wang et al. (2018) propose an unsupervised data augmentation method for NMT by replacing words in both source and target sentences based on hamming distance. Gao et al. (2019) propose a method that replaces words with a weighted combination of semantically similar words. Recently, Nguyen et al. (2020) propose an in-domain augmentation method by diversifying the available bitext data using multiple forward and backward models. In their follow-up work (Nguyen et al., 2021), they extend the idea to unsupervised MT (UMT) using a cross-model distillation method, where one UMT model’s synthetic output is used as input for another UMT model.

Summary Most of the previous work on improving BT involve either training iteratively or combining BT with self-training using monolingual data blindly without noticing the distributional differences between the monolingual and bitext data. In contrast, in AUGVIC we systematically parameterize the generation of new training samples from the original parallel data. Moreover, the combination of our augmented vicinal samples with monolingual data makes the NMT models more robust and attenuates the prevailing distributional gap.

3 Method

Let s and t denote the source and target languages respectively, and $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ denote the bitext training corpus containing N sentence pairs with x_i and y_i coming from s and t languages, respectively. Also, let $\mathcal{M}_{s \rightarrow t}$ is an NMT model that can translate sentences from s to t , and $\mathcal{D}_{\text{mono}}^t = \{y_j\}_{j=1}^M$ denote the monolingual corpus in the target language t containing M sentences.

3.1 Traditional Back-Translation

Traditional back-translation (Sennrich et al., 2016) leverages the target-side monolingual corpus. With the aim to train a source-to-target model $\mathcal{M}_{s \rightarrow t}$, it first trains a reverse intermediate model $\mathcal{M}_{t \rightarrow s}$ using the given bitext \mathcal{D} , and use it to translate the extra target-side monolingual data $\mathcal{D}_{\text{mono}}^t$ into source language. This yields a synthetic bitext corpus $\mathcal{D}_{\text{syn}} = \{\mathcal{M}_{t \rightarrow s}(y_j), y_j\}_{j=1}^M$. Then a final model $\mathcal{M}_{s \rightarrow t}$ is trained on $\{\mathcal{D} \cup \mathcal{D}_{\text{syn}}\}$ usually by upsampling \mathcal{D} to keep the original and synthetic bitext pairs to a certain ratio (generally 1:1).

3.2 AUGVIC: Exploiting Bitext Vicinity

For low-resource languages, the amount of available parallel data is limited, hindering training of a good MT system. Moreover, the target language pairs can be quite different (e.g., morphologically, topic distribution) from the high-resource ones, making the translation task more difficult (Chen et al., 2019). Also, acquiring large and relevant monolingual corpora in the target language is difficult in low-resource settings and can be quite expensive. The domain mismatch between the monolingual and bitext data is another issue with the traditional back-translation as mentioned in §1.

With the aim to improve model generalization, the core idea of AUGVIC is to leverage the *vicinal*

samples of the given bitext rather than using extra monolingual data. The addition of bitext vicinity also alleviates the domain mismatch issue since the augmented data distribution does not change much from the original bitext distribution. Figure 1 shows an illustrative example of AUGVIC, which works in three basic steps to train a model:

- (i) Generate vicinal samples \tilde{y}_i of the target sentences (y_i) in the bitext data \mathcal{D} .
- (ii) Produce source-side translations \tilde{x}_i of the vicinal samples to generate synthetic bitext $\tilde{\mathcal{D}}$.
- (iii) Train the final source-to-target MT model $\mathcal{M}_{s \rightarrow t}$ using $\{\mathcal{D} \cup \tilde{\mathcal{D}}\}$.

AUGVIC, however, is not mutually exclusive to the traditional back-translation and can be used together when relevant monolingual data is available. In the following, we describe how each of these steps are operationalized with NMT models.

3.2.1 Generation of Vicinal Samples

We first generate vicinal samples for each eligible target sentence y_i in the bitext $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$. Let $\mathcal{V}(\tilde{y}_i|y_i)$ denote the vicinity distribution around y_i , we create a corpus of vicinal samples as:

$$\tilde{y}_i \sim \mathcal{V}(\tilde{y}_i|y_i) \quad (1)$$

We generate vicinal samples for sentences having lengths between 3 and 100, and \mathcal{V} can be modeled with existing syntactic and semantic alternation methods like language model (LM) augmentation (Kobayashi, 2018; Wu et al., 2018; Shi et al., 2020; Bari et al., 2021), paraphrase generation (Li et al., 2018), constrained summarization (Laban et al., 2020), and similar sentence retrieval (Du et al., 2020). Most of these methods are supervised requiring extra annotations. Instead, in AUGVIC, we adopt an unsupervised LM augmentation, which makes the framework more robust and flexible to use. Specifically, we use a pretrained XLM-R masked LM (Conneau et al., 2020a) parameterized by θ_{xlmr} as our vicinal model. Thus, the vicinity distribution is defined as $\mathcal{V}(\tilde{y}_i|y_i, \theta_{\text{xlmr}})$.

Note that we treat the vicinal model as an external entity, which is not trained/fine-tuned. This disjoint characteristic gives our framework the flexibility to replace θ_{xlmr} even with a better monolingual LM for a specific target language, which in turn makes AUGVIC extendable to utilize stronger LMs that may come in the future.

In a masked LM, one can mask out a token at any position and ask the model to predict at that position. For a meaningful and informed augmentation, we mask out the tokens *successively* (one at a time) up to a required number determined by a diversity ratio, $\rho \in (0, 1)$. For a sentence of length ℓ , the successive augmentation can generate at most $(2^\ell - 1) \times k$ vicinal samples, where k is the number of output tokens chosen for each masked position. We use $k = 1$, and pick the one with the highest probability ensuring that it does not match the original token at the masked position. The diversity ratio (ρ) controls how much diverse the vicinal samples can be from the original sentence, and is selected using one of the following two ways:

- **Fixed diversity ratio** Here we use a fixed value for ρ , and select $t = \ell \times \rho$ tokens to mask out. We then generate new vicinity samples by predicting new tokens in those masked positions.
- **Dynamic diversity ratio** Instead of using a fixed value, in this approach we set the diversity ratio dynamically taking the sentence length into consideration. This allows finer-level control for diversification — the longer the sentence is, the smaller should its diversification ratio be. The intuition is that for long sentences, a larger value of ρ will produce vicinal samples which will be far away from the original sample. Specifically, we use the following piece-wise function to find the number of tokens to mask out dynamically:

$$t = \begin{cases} \max(\ell \times a, t_{\min}) & ; \text{if } \ell \leq 20 \\ \min(\frac{\ell}{h} \times b, t_{\max}) & ; \text{otherwise} \end{cases} \quad (2)$$

where t_{\min} and t_{\max} are hyperparameters and represent the minimum and maximum number of tokens to be replaced by the masked LM. The other hyperparameters a , b , and h play the same role as the diversity ratio ρ .

Since we predict tokens for replacement one at a time, we can make the prediction in any of the permutation order of t . So, the maximum number of possible augmentation for a sentence of length ℓ is $\gamma = \binom{\ell}{t} \times t!$. We perform *stochastic sampling* from the distribution of γ to select N' vicinal samples. We have added an analysis on the effect of diversity ratio ρ in AUGVIC in §5.5.

3.2.2 Generation of Synthetic Bitext Data

Our objective is to train a source-to-target MT model $\mathcal{M}_{s \rightarrow t}$. So far, we have the bitext $\mathcal{D} =$

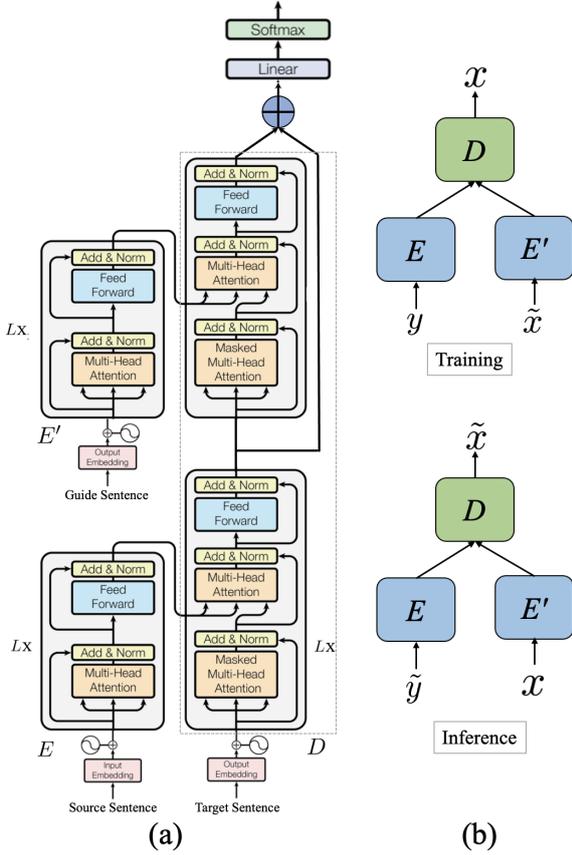


Figure 2: (a) Our proposed model for guided back-translation; (b) its training and inference method.

$\{(x_i, y_i)\}_{i=1}^N$ and target-side monolingual data $\tilde{\mathcal{D}}^t = \{\tilde{y}_j\}_{j=1}^{N'}$ which are vicinal to the original target in \mathcal{D} . We need a reverse intermediate target-to-source MT model $\mathcal{M}_{t \rightarrow s}$ to translate \tilde{y}_j into \tilde{x}_j , which will give us the synthetic bitext data $\tilde{\mathcal{D}}$. For this, we experiment with two different models.

(a) Pure Back-Translation (PBT) This is similar to back-translation (§3.1), where we first train the reverse MT model $\mathcal{M}_{t \rightarrow s}$ using the given bitext \mathcal{D} . We then use $\mathcal{M}_{t \rightarrow s}$ to translate the target-side vicinal samples $\tilde{y}_j \sim \tilde{\mathcal{D}}^t$ into \tilde{x}_j . This gives a synthetic bitext $\tilde{\mathcal{D}} = \{(\tilde{x}_j, \tilde{y}_j)\}_{j=1}^{N'}$. We use the Transformer architecture (Vaswani et al., 2017) as our reverse intermediate NMT model $\mathcal{M}_{t \rightarrow s}$.

(b) Guided Back-Translation (GBT) In the illustrative example (Figure 1), we can identify three kinds of pairs: (i) the bitext (x_i, y_i) , (ii) the vicinal (y_i, \tilde{y}_i) , and (iii) the synthetic pair $(\tilde{x}_i, \tilde{y}_i)$. Here, y_i is the original translation of source sentence x_i and \tilde{y}_i is the vicinal sample, which can be seen as a perturbation of y_i . Hence, we can assume that \tilde{x}_i will also be similar to (perturbed) x_i . Our goal is to leverage this extra relational knowledge to improve

the translation quality of \tilde{x}_i when generating the synthetic bitext $\tilde{\mathcal{D}}$. Specifically, we use the original source x_i as a guide for generating the synthetic translation \tilde{x}_i of the target-side vicinal sample \tilde{y}_i .

$$\tilde{x}_i = \mathcal{M}_{t \rightarrow s}(\tilde{y}_i | x_i) \quad (3)$$

For this, we propose a model based on the Transformer architecture which has two encoders - one for the source sentence (E) and another for the guide sentence (E'), and a decoder (D) (Figure 2). We use the same architecture with the exception that now we have two identical encoders (E and E'). Both the encoders have a stack of L layers, while the decoder has $(L + 1)$ layers.

Training & Inference: We train this model with a dataset of triplets containing (y, \tilde{x}, x) , where (x, y) comes from the original bitext and \tilde{x} is a vicinal sample of x to guide the decoder in generating x . Each of the first L layers of the decoder performs cross-attention on $E(y)$ resulting in decoder states $D^{(L)}(x_{<t}|y)$ at time step t , while the final decoder layer attends on $E'(\tilde{x})$ resulting in a second set of decoder states $D^{(L+1)}(x_{<t}|y, \tilde{x})$. The two sets of decoder states are then interpolated by taking a convex combination before passing it to a linear layer followed by the Softmax token prediction.

$$\lambda D^{(L)}(x_{<t}|y) + (1 - \lambda) D^{(L+1)}(x_{<t}|y, \tilde{x}) \quad (4)$$

where λ is a hyperparameter that controls the relative contributions from the two encoders, $E(y)$ and $E'(\tilde{x})$, in generating x by the decoder D .

To generate the synthetic bitext $\tilde{\mathcal{D}}$, we need to translate \tilde{y} , which will be guided by x . So during *inference*, we feed \tilde{y} to E and x to E' to autoregressively generate \tilde{x} with beam search decoding.

3.2.3 Training of the Final Model

We combine the original bitext \mathcal{D} and the synthetic bitext $\tilde{\mathcal{D}}$ generated from the previous step to train our final source-to-target model $\mathcal{M}_{s \rightarrow t}$. We use the standard Transformer as our final model.

4 Experimental Setup

4.1 Datasets and Evaluation Metrics

We conduct experiments on four low-resource language pairs: English (En) to/from Bangla (Bn), Tamil (Ta), Nepalese (Ne), and Sinhala (Si). Table 1 presents the source of the collected datasets and their domains for each language pair.

Pair	Data-Source	Train & Dev	Test
En-Bn	Islam (2018)	Mixed	Mixed
En-Ta	Ramasamy et al. (2014)	News, Bible, Cinema	News, Bible, Cinema
En-Ne	Guzmán et al. (2019)	Bible, GV, PTB, Ubuntu	Wikipedia
En-Si	Guzmán et al. (2019)	Opens subtitles, Ubuntu	Wikipedia

Table 1: Sources and domains of the datasets.

Even though the En-Bn dataset size is relatively small ($\sim 72\text{K}$ pairs), the quality of the bitext is rich, and it covers a diverse set of domains including literature, journalistic texts, instructive texts, administrative texts, and texts treating external communication. Here the distributions in train and test splits are about the same. For En-Ta, the train and test domains are similar, mostly coming from the news ($\sim 66.43\%$). For En-Ne and En-Si, we use the datasets from (Guzmán et al., 2019), where the train and test domains are different. Although these two datasets are comparatively larger ($\sim 600\text{K}$ pairs each), the quality of the bitext is poor, requiring further cleaning and deduplication.

Table 2 presents the dataset statistics after deduplication where the last column specifies the number of augmented data by our method AUGVIC (§3.2.1). For a fair comparison with the traditional back-translation, we experiment with the same amount of target-side monolingual data from three domains: news, wiki, and gnome. We collected and cleaned News, Wiki, and Gnome datasets from News-crawl, Wiki-dumps, and Gnome localization guide, respectively. For some languages, the amount of specific domain monolingual data is limited, where we added additional monolingual data of that language from Common Crawl.

Following previous work (Guzmán et al., 2019; Nguyen et al., 2020), we report the tokenized BLEU (Papineni et al., 2002) when translating from English to other languages, and detokenized SacreBLEU (Post, 2018) when translating from other languages to English for all our experiments.

Pair	Train	Dev	Test	Augmented (AUGVIC/Mono)
En-Bn	70,854	500	500	$\approx 460\text{K}$
En-Ta	166,851	1000	2000	$\approx 1300\text{K}$
En-Ne	234,514	2559	2835	$\approx 1500\text{K}$
En-Si	571,213	2898	2766	$\approx 1500\text{K}$

Table 2: Dataset statistics after deduplication.

4.2 Baselines

We compare AUGVIC with the following baselines:

(i) **Bitext baseline** is the model trained with the bitext given with the dataset.

(ii) **Upsample baseline** Here we upsample the bitext to the same amount of AUGVIC’s data.

(iii) **Diversification baseline** Nguyen et al. (2020) diversifies the original parallel data by using the predictions of multiple forward and backward NMT models. Then they merge the augmented data with the original bitext on which the final NMT model is trained. Their method is directly comparable to AUGVIC, as both methods diversify the original bitext, but in different ways.

4.3 Model Settings

We use the Transformer (Vaswani et al., 2017) implementation in Fairseq (Ott et al., 2019). We follow the basic architectural settings from (Guzmán et al., 2019), which establishes some standards for low-resource MT. For low-resource “Bitext baseline”, they use a smaller (5-layer) Transformer architecture as the dataset is small, while for larger datasets (*e.g.*, with additional synthetic data) they use a bigger (6-layer) model.² To keep the architecture the same in the respective rows (Table 3), we use a 6-layer model for “Upsample baseline” and 5-layer for “Bitext baseline”. More specifically, for datasets with less than a million bitext pairs, we use an architecture with 5 encoder and 5 decoder layers, where the number of attention heads, embedding dimension, and inner-layer dimension are respectively 8, 512, and 2048. Otherwise, we use a larger Transformer architecture with 6 encoder and 6 decoder layers with the number of attention heads, embedding dimension, and inner-layer dimension of 16, 1024, and 4096, respectively.

After deduplication, we tokenize non-English data using the Indic NLP Library.³ We use the sentencepiece library⁴ to learn the joint Byte-Pair-Encoding (BPE) of size 5000 symbols for each of

²<https://github.com/facebookresearch/flores/>

³https://github.com/anoopkunchukuttan/indic_nlp_library

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

Setting	Data	En-Bn		En-Ta		En-Ne		En-Si	
		→	←	→	←	→	←	→	←
Baseline	Bitext × Upsample	13.21	21.18	11.58	26.29	4.59	8.34	1.96	7.45
		16.59	25.51	12.15	27.71	4.16	7.79	1.81	6.93
Diversification	+ Nguyen et al. (2020)	17.54	26.11	12.74	28.54	5.7	8.9	2.2	8.2
	+ AUGVIC	18.03	26.96	12.93	28.68	6.47	10.65	3.66	9.27
Extra mono. data	+ BT-Mono (News)	18.81	27.11	13.51	29.38	6.44	12.48	3.56	11.75
	+ BT-Mono (Wiki)	18.52	26.33	13.23	29.01	6.91	13.02	3.91	11.86
	+ AUGVIC+ BT-Mono (News)	19.98	28.14	13.87	30.15	6.80	13.12	4.94	11.89
	+ AUGVIC+ BT-Mono (Wiki)	20.39	28.48	13.89	30.14	7.27	13.52	5.24	12.09

Table 3: Detokenized Sacre-BLEU scores for {Bn, Ta, Ne, Si} → En and tokenized BLEU fro En → {Bn, Ta, Ne, Si}. “BT-Mono” stands for traditional back-translation with extra target-side monolingual data (§3.1).

the language pair over the raw English and tokenized non-English bitext training data.

We tuned the hyper-parameters a , b , h , t_{min} , t_{max} in Eq. 2 and λ in Eq. 4 by small-scale experiments on the validation-sets. We found $a = 0.5$, $b = 2.5$, $h = 10$, $t_{min} = 1$, and $t_{max} = 20$ work better. We tuned λ within the range of 0.5 to 0.9. In general, we observe that for smaller sentences (length ≤ 20), 50-60% successive-token-replacement works better while for longer sentences (length > 20), 20-30% token-replacement performs better.

Following Guzmán et al. (2019), we train all the models upto a maximum epoch of 100 with early-stopping enabled based on the validation loss. We use the beam-search-decoding for inference. All the reported results for AUGVIC use dynamic diversity ratio for generating vicinal samples unless otherwise specified.

5 Results and Analysis

In this section, we present our results and the analysis of our proposed methods.

5.1 Comparison with Bitext & Diversification

Table 3 presents the BLEU scores on the eight translation tasks. First, we compare our model AUGVIC with the model trained on the original parallel data (Bitext). AUGVIC consistently improves the performance over all the tested language pairs, gaining about +2.76 BLEU scores on average. Specifically, AUGVIC achieves the absolute improvements of 4.28, 5.78, 1.35, 2.39, 1.88, 2.31, 1.70, and 1.82 over the Bitext for En-Bn, Bn-En, En-Ta, Ta-En, En-Ne, Ne-En, En-Si, and Si-En, respectively.

For a fair comparison, in another experiment, we upsample the bitext data to make it similar to the amount of AUGVIC’s data. From the *Upsample* re-

sults (with a 6-layer architecture) reported in Table 3, we see that even though it increases the BLEU scores for En to/from {Bn, Ta}, it has negative impacts on En to/from {Ne, Si} where it degrades the performance. Overall, AUGVIC achieves 1.75 BLEU score improvements on an average over the Upsample baseline.

The comparison with the diversification strategy proposed by Nguyen et al. (2020) reveals that AUGVIC outperforms their method by 0.84 BLEU scores on average. To be specific, our method gets 0.49, 0.85, 0.19, 0.14, 0.77, 1.75, 1.46, and 1.07 absolute BLEU improvements over their approach for En-Bn, Bn-En, En-Ta, Ta-En, En-Ne, Ne-En, En-Si, and Si-En, respectively.

The data diversification method of Nguyen et al. (2020) relies heavily on the performance of base models (Bitext). From Table 3, we see that the performance of base models are poor for En to/from {Ne, Si}, which impacts their augmented data generation process (diversification). However, the better performance of AUGVIC in those languages indicates that vicinal samples generated in our method are more diverse with better quality and less prone to the noise in base models.

5.2 Vicinal Samples with Extra Relevant Monolingual Data

We further explore the performance of AUGVIC by experimenting with the traditional back-translation method (§3.1) using the same amount of monolingual data. To perceive the variability, we choose to experiment with extra monolingual data from two *relevant* but different sources - newscrawl (BT-Mono (News)) and Wikipedia (BT-Mono (Wiki)). From the results in Table 3, we see that standard back-translation improves the scores in both cases, proving that extra relevant monolingual data helps

Interm. BT system	En-Bn		En-Ta		En-Ne		En-Si	
	→	←	→	←	→	←	→	←
Pure BT	18.03	26.96	12.93	28.68	6.47	10.65	3.66	9.27
Guided BT	18.18	27.35	13.17	29.05	4.81	8.62	2.16	7.71

Table 4: Comparison between two **intermediate** reverse back-translation (**BT**) systems in AUGVIC.

for low-resource MT significantly.

To understand the exclusivity of the vicinal samples of AUGVIC from the external related monolingual data, we perform another set of experiments where we added both the AUGVIC’s augmented data with the extra monolingual data and trained along with the Bitext data. From Table 3, we see that the combination of datasets improves the BLEU scores by 1.02 and 0.73 on average on the two relevant data sources (News and Wiki). From this, we can conclude that vicinal samples of AUGVIC make the NMT models more robust in the presence of the relevant monolingual data and can be used together when available.

5.3 Pure vs. Guided: Which One is Better?

For all the results of AUGVIC presented in Table 3, we use the pure back-translation (BT) method (§3.2.2(a)) as the reverse intermediate model. We compare the performance of the guided BT (§3.2.2(b)) with the pure BT method as the reverse intermediate model in Table 4. From the results, we observe that the guided BT achieves better results in $En \leftrightarrow \{Bn, Ta\}$, while the pure BT achieves better in $En \leftrightarrow \{Ne, Si\}$ translation tasks.

We investigated why the guided BT performed poorly in $En \leftrightarrow \{Ne, Si\}$ tasks, and found that compared to the En-Bn and En-Ta bitexts, the original bitexts of En-Ne and En-Si languages are very noisy (e.g., bad sentence segmentation, code-mix data), which propagates further noise while using the target translation as a guide for translating the vicinal samples. The diminishing results while up-sampling in these two languages (Table 3) supports this claim. From these results, we can say that the better the original bitext quality is, the better the synthetic bitext will be for the guided BT.

5.4 AUGVIC with Relevant and Distant-domain Monolingual Data

To verify how traditional back-translation and AUGVIC perform with with monolingual data from related vs. distant domains, we perform another set of experiments on En to/from $\{Bn, Ta\}$. For both the language pairs (§4.1), *News* can (roughly) be

BT-mono Domain	Data	En-Bn		En-Ta	
		→	←	→	←
	Bitext	13.21	21.18	11.58	26.29
<i>News</i> (relevant)	+ BT	18.81	27.11	13.51	29.38
	+ AUGVIC+ BT	19.98	28.14	13.87	30.15
<i>gnome</i> (distant)	+ BT	17.14	26.05	12.55	27.91
	+ AUGVIC+ BT	18.86	27.56	13.59	29.89

Table 5: Effect of relevant and distant domain monolingual data in back-translation with AUGVIC. We use *News* as “relevant” and *gnome* as “distant” domain.

considered as relevant compared to *gnome*,⁵ which can be considered as distant domain. We use *pure BT* as the intermediate reverse back-translation system for generating synthetic data in AUGVIC in this set of experiments.

From Table 5, we see that traditional back-translation (+ BT) improves the BLEU scores over the Bitext by 4.14 and 2.85 on average for relevant- and distant-domain monolingual data, respectively, yielding higher gains for relevant domain, as expected. The addition of vicinal data by AUGVIC (+ AUGVIC+ BT) further improves the scores in both cases; interestingly, the relative improvements are higher in the distant-domain case. Specifically, the average BLEU score improvements over Bitext for relevant- and distant-domain data with AUGVIC+BT are 4.97 and 4.41, respectively. Comparing this with BT only, the BLEU score difference between relevant and distant domains has been reduced from 1.29 to 0.56. This indicates that AUGVIC helps to bridge the domain gap between relevant and distant-domain distributions in traditional BT with monolingual data.

In principle, for vicinal samples, the synthetic-pair generation capability of the reverse intermediate target-to-source MT model should be better than generating from an arbitrary monolingual data as it could be a distant distribution compared to the bitext. Judging by the amount of diverse data used for training the language model, we can safely assume that it is a diverse knowledge source (Conneau et al., 2020b) compared to the training bitext samples. Data that performs well on the reverse

⁵<http://opus.nlpl.eu/GNOME.php>

intermediate target-to-source MT system can be extrapolated from the knowledge-base as vicinal-distribution with the controlled diversity ratio function (Eq. 2). Moreover, to achieve more diversity, the use of multiple different language models is also compatible in AUGVIC.

5.5 Effect of Diversity Ratio in AUGVIC

For monolingual data, it could be challenging to identify domain discrepancy with the training/testing bitext data, and there is no parameter in the traditional BT method to control this distributional mismatch. However, in AUGVIC we can control the distributional drift of the generated vicinal samples from the original training distribution by varying the diversity ratio ρ .

Theoretically, it is possible to sample the same distribution using dynamic and static diversity. However, dynamic diversity is more flexible to perform hyperparameter-tuning and to prevent potential outliers. The term l/h in Eq. 2 represents pseudo-segmentation (h segments) of a large sentence of length l , and b represents the same intuition as ρ . Apart from these, t_{min} and t_{max} prevents irregular-samples: (i) t_{min} ensures that there should be at least some changes in the augmented sample, (ii) t_{max} makes sure that the generated-samples from LM do not diverge too much from the vicinity.

To understand the effect of the diversity ratio in AUGVIC, we perform another set of experiments. We choose to use En to/from {Bn, Ne} for this experiments, where we selected at most two vicinal samples from each of the target sentence in original bitext. We investigate the effect of both *dynamic* and *fixed* diversity ratio in AUGVIC’s vicinal sample generation (§3.2.1). For fixed diversity ratio we use ρ values 0.1, 0.3, 0.5, and 0.8, while for dynamic diversity ratio we use $a = 0.5, b = 2.5$, and $h = 10$ for controlling the diversity.

We present these experimental results in Table 6, from where we see that the dynamic diversity ratio performs better in three out of four tasks. For the fixed diversity ratio, we see the variation in results for different values of ρ . In all the four tasks, the diversity ratio $\rho = 0.8$ gives the least scores. On average, we get the better results with $\rho = \{0.3, 0.5\}$. These experiments suggest that higher diversity values may induce noise and lower diversity values may not diversify the data enough to benefit the final NMT model.

AUGVIC diversity ratio	En-Bn		En-Ne	
	→	←	→	←
Dynamic	17.69	26.61	6.21	10.25
Fixed				
$\rho = 0.1$	17.34	25.98	5.98	10.03
$\rho = 0.3$	17.52	26.19	6.19	10.36
$\rho = 0.5$	17.48	26.49	6.05	10.38
$\rho = 0.8$	17.19	25.01	5.82	9.89

Table 6: Effect of diversity ratio ρ while generating vicinal samples in AUGVIC (§3.2.1).

6 Conclusion

We have presented an in-domain data augmentation framework AUGVIC by exploiting the bitext vicinity for low-resource NMT. Our method generates vicinal samples by diversifying sentences of the target language in the bitext in a novel way. It is simple yet effective and can be quite useful when extra in-domain monolingual data is limited.

Extensive experiments with four low-resource language pairs comprising data from different domains show the efficacy of AUGVIC. Our method is not only comparable with traditional back-translation with in-domain monolingual data, it also makes the NMT models more robust in the presence of relevant monolingual data. Moreover, it bridges the distributional gap for out-of-domain monolingual data when using together.

References

- Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George F. Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. [Massively multilingual neural machine translation in the wild: Findings and challenges](#). *CoRR*, abs/1907.05019.
- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2018. A robust self-learning method for fully unsupervised cross-lingual mappings of word embeddings. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 789–798.
- M Saiful Bari, Tasnim Mohiuddin, and Shafiq Joty. 2021. Uxla: A robust unsupervised data augmentation framework for cross-lingual nlp. In *Proceedings of The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (ACL-IJCNLP 2021)*, Online. Association for Computational Linguistics.

- Isaac Caswell, Ciprian Chelba, and David Grangier. 2019. [Tagged back-translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 53–63, Florence, Italy. Association for Computational Linguistics.
- Olivier Chapelle, Jason Weston, Léon Bottou, and Vladimir Vapnik. 2001. [Vicinal risk minimization](#). In T. K. Leen, T. G. Dietterich, and V. Tresp, editors, *Advances in Neural Information Processing Systems 13*, pages 416–422. MIT Press.
- Peng-Jen Chen, Jiajun Shen, Matthew Le, Vishrav Chaudhary, Ahmed El-Kishky, Guillaume Wenzek, Myle Ott, and Marc’Aurelio Ranzato. 2019. [Facebook AI’s WAT19 Myanmar-English translation task submission](#). In *Proceedings of the 6th Workshop on Asian Translation*, pages 112–122, Hong Kong, China. Association for Computational Linguistics.
- Yong Cheng, Qian Yang, Yang Liu, Maosong Sun, and Wei Xu. 2017. [Joint training for pivot-based neural machine translation](#). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 3974–3980.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020a. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020b. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. [Word translation without parallel data](#). *CoRR*, abs/1710.04087.
- Raj Dabre, Chenhui Chu, and Anoop Kunchukuttan. 2020. [A survey of multilingual neural machine translation](#). *ACM Comput. Surv.*, 53(5).
- Jingfei Du, Edouard Grave, Beliz Gunel, Vishrav Chaudhary, Onur Celebi, Michael Auli, Ves Stoyanov, and Alexis Conneau. 2020. [Self-training improves pre-training for natural language understanding](#).
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- Sergey Edunov, Myle Ott, Marc’Aurelio Ranzato, and Michael Auli. 2020. [On the evaluation of machine translation systems trained with back-translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2836–2846, Online. Association for Computational Linguistics.
- Marzieh Fadaee, Arianna Bisazza, and Christof Monz. 2017. [Data augmentation for low-resource neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 567–573, Vancouver, Canada. Association for Computational Linguistics.
- Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016a. [Multi-way, multilingual neural machine translation with a shared attention mechanism](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 866–875, San Diego, California. Association for Computational Linguistics.
- Orhan Firat, Baskaran Sankaran, Yaser Al-onazian, Fatos T. Yarman Vural, and Kyunghyun Cho. 2016b. [Zero-resource translation with multi-lingual neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 268–277, Austin, Texas. Association for Computational Linguistics.
- Fei Gao, Jinhua Zhu, Lijun Wu, Yingce Xia, Tao Qin, Xueqi Cheng, Wengang Zhou, and Tie-Yan Liu. 2019. [Soft contextual data augmentation for neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5539–5544, Florence, Italy. Association for Computational Linguistics.
- Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor O. K. Li. 2018. [Universal neural machine translation for extremely low resource languages](#). *CoRR*, abs/1802.05368.
- Francisco Guzmán, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. [The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6098–6111, Hong Kong, China. Association for Computational Linguistics.
- Hany Hassan, Anthony Aue, C. Chen, Vishal Choudhary, 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.
- Junxian He, Jiatao Gu, Jiajun Shen, and Marc’Aurelio Ranzato. 2020. [Revisiting self-training for neural sequence generation](#). In *International Conference on Learning Representations*.
- Vu Cong Duy Hoang, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. [Iterative back-translation for neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 18–24, Melbourne, Australia. Association for Computational Linguistics.
- M. A. A. Mumin; M. H. Seddiqui; M. Z. Iqbal; M. J. Islam. 2018. [Supara0.8m: A balanced english-bangla parallel corpus](#).
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Yunsu Kim, Petre Petrov, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2019. [Pivot-based transfer learning for neural machine translation between non-English languages](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 866–876, Hong Kong, China. Association for Computational Linguistics.
- Sosuke Kobayashi. 2018. [Contextual augmentation: Data augmentation by words with paradigmatic relations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 452–457, New Orleans, Louisiana. Association for Computational Linguistics.
- Tom Kocmi and Ondřej Bojar. 2018. [Trivial transfer learning for low-resource neural machine translation](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 244–252, Brussels, Belgium. Association for Computational Linguistics.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39, Vancouver. Association for Computational Linguistics.
- Philippe Laban, Andrew Hsi, John Canny, and Marti A. Hearst. 2020. [The summary loop: Learning to write abstractive summaries without examples](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, ACL 2020, Online, July 5-10, 2020*, pages 5135–5150. Association for Computational Linguistics.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2021. [{GS}hard: Scaling giant models with conditional computation and automatic sharding](#). In *International Conference on Learning Representations*.
- Zichao Li, Xin Jiang, Lifeng Shang, and Hang Li. 2018. [Paraphrase generation with deep reinforcement learning](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3865–3878, Brussels, Belgium. Association for Computational Linguistics.
- Yinhan Liu, Jiatao Gu, Naman Goyal, Xian Li, Sergey Edunov, Marjan Ghazvininejad, Mike Lewis, and Luke Zettlemoyer. 2020. [Multilingual denoising pre-training for neural machine translation](#).
- Tasnim Mohiuddin, M Saiful Bari, and Shafiq Joty. 2020. [LNMap: Departures from isomorphic assumption in bilingual lexicon induction through non-linear mapping in latent space](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2712–2723, Online. Association for Computational Linguistics.
- Tasnim Mohiuddin and Shafiq Joty. 2019. [Revisiting adversarial autoencoder for unsupervised word translation with cycle consistency and improved training](#). 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 3857–3867, Minneapolis, Minnesota. Association for Computational Linguistics.
- Tasnim Mohiuddin and Shafiq Joty. 2020. [Unsupervised Word Translation with Adversarial Autoencoder](#). *Computational Linguistics*, 46(2):257–288.
- Graham Neubig and Junjie Hu. 2018. [Rapid adaptation of neural machine translation to new languages](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 875–880, Brussels, Belgium. Association for Computational Linguistics.
- Xuan-Phi Nguyen, Shafiq Joty, Thanh-Tung Nguyen, Wu Kui, and Ai Ti Aw. 2021. [Cross-model Back-translated Distillation for Unsupervised Machine Translation](#). In *Thirty-eighth International Conference on Machine Learning, ICML’21*, Virtual.
- Xuan-Phi Nguyen, Shafiq R. Joty, Kui Wu, and Ai Ti Aw. 2020. [Data diversification: A simple strategy](#).

- for neural machine translation. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- 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.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. 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.
- Loganathan Ramasamy, Ondřej Bojar, and Zdeněk Žabokrtský. 2014. [EnTam: An english-tamil parallel corpus \(EnTam v2.0\)](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich and Biao Zhang. 2019. [Revisiting low-resource neural machine translation: A case study](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 211–221, Florence, Italy. Association for Computational Linguistics.
- Linqing Shi, Danyang Liu, Gongshen Liu, and Kui Meng. 2020. Aug-bert: An efficient data augmentation algorithm for text classification. In *Communications, Signal Processing, and Systems*, pages 2191–2198, Singapore. Springer Singapore.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC’12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- 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.
- Xinyi Wang, Hieu Pham, Zihang Dai, and Graham Neubig. 2018. [SwitchOut: an efficient data augmentation algorithm for neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 856–861, Brussels, Belgium. Association for Computational Linguistics.
- Lijun Wu, Yiren Wang, Yingce Xia, Tao Qin, Jianhuang Lai, and Tie-Yan Liu. 2019. [Exploiting monolingual data at scale for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4207–4216, Hong Kong, China. Association for Computational Linguistics.
- Xing Wu, Shangwen Lv, Liangjun Zang, Jizhong Han, and Songlin Hu. 2018. [Conditional BERT contextual augmentation](#). *CoRR*, abs/1812.06705.
- Ziang Xie, Sida I. Wang, Jiwei Li, Daniel Lévy, Aiming Nie, Dan Jurafsky, and Andrew Y. Ng. 2017. Data noising as smoothing in neural network language models. *CoRR*, abs/1703.02573.
- Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tieyan Liu. 2020. [Incorporating bert into neural machine translation](#). In *International Conference on Learning Representations*.
- 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.