Tag-less Back-Translation

Idris Abdulmumin (✉ iabdulmumin@abu.edu.ng)

Ahmadu Bello University <https://orcid.org/0000-0002-3795-8381>

Bashir Shehu Galadanci

Bayero University

Aliyu Garba

Ahmadu Bello University

Research Article

Keywords: tagged back-translation, tag-less back-translation, neural machine translation, machine translation, natural language processing

Posted Date: April 30th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-465941/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Machine Translation on November 2nd, 2021. See the published version at <https://doi.org/10.1007/s10590-021-09284-y>.

Tag-less Back-Translation

Idris Abdulmumin · Bashir Shehu
Galadanci · Aliyu Garba

This work is supported by the National Information Technology Development Agency under the National Information Technology Development Fund PhD Scholarship Scheme 2018

Received: date / Accepted: date

Abstract An effective method to generate a large number of parallel sentences for training improved neural machine translation (NMT) systems is the use of the back-translations of the target-side monolingual data. The standard back-translation method has been shown to be unable to efficiently utilize the available huge amount of existing monolingual data because of the inability of translation models to differentiate between the authentic and synthetic parallel data during training. Tagging, or using gates, has been used to enable translation models to distinguish between synthetic and authentic data, improving standard back-translation and also enabling the use of iterative back-translation on language pairs that underperformed using standard back-translation. In this work, we approach back-translation as a domain adaptation problem, eliminating the need for explicit tagging. In the approach – *tag-less back-translation* – the synthetic and authentic parallel data are treated as out-of-domain and in-domain data respectively and, through pre-training and fine-tuning, the translation model is shown to be able to learn more efficiently from them during training. Experimental results have shown that the approach outperforms the standard and tagged back-translation approaches on low resource English-Vietnamese and English-German neural machine translation.

Keywords tagged back-translation · tag-less back-translation · neural machine translation · machine translation · natural language processing

Idris Abdulmumin
Computer Science, Ahmadu Bello University, Zaria, Kaduna, Nigeria.
Tel.: +234-806-295-5509
E-mail: iabdulmumin@abu.edu.ng

Bashir Shehu Galadanci
Software Engineering, Bayero University, Kano, Kano, Nigeria.
E-mail: bsgaladanci.se@buk.edu.ng

Aliyu Garba
Computer Science, Ahmadu Bello University, Zaria, Kaduna, Nigeria.
E-mail: algarba@abu.edu.ng

1 Introduction

Neural Machine Translation (NMT) [2, 18, 48] has been the state-of-the-art approach for machine translation in recent years [16, 38], outperforming Phrase-Based Statistical Machine Translation [33] when qualitative parallel data between the languages is available in abundance [55]. This training dataset is usually scarce and expensive to compile for many language pairs. Recently, researchers have proposed methods to exploit the easier-to-get monolingual data of one or both of the languages to augment the available parallel data and improve the performance of the translation models. Such methods include integrating a language model [20], back-translation [45, 22, 19], forward translation [53] and dual learning [21]. The back-translation approach is simple and has been the most effective technique yet for NMT [16, 22]. The method involves training a target-to-source (backward) model on the available authentic bitext. The backward model is then used to translate a large amount of monolingual sentences in the target language into synthetic source sentences, generating the synthetic parallel data. The authentic and synthetic parallel data are then mixed to train a source-to-target (forward) model.

It has been shown that as the amount of monolingual data used in back-translation continues to increase, a point is reached when the model stops learning useful representation and, therefore, the performance of the model starts to drop. This is because the usually noise-infested synthetic data starts to overwhelm the authentic data and the model starts to completely unlearn the correct parameters it learns from the authentic training data [17]. Extensive studies by [16] have shown that in low resource NMT, noising beam search outputs improve the models more than other generation methods such as sampling. The authors claimed that the method enhances source-side diversity. But the works of [6, 52] found that the noising technique is only a form of tagging, indicating to the model that the noised data is back-translated, enabling it to treat the synthetic data as belonging to a different domain. The model then learns different representations, optimally, from the two data. They, instead, introduced the use of explicit tags (and gates) to indicate synthetic inputs. The tagging approach was shown to outperform the standard back-translation.

In this work, we approach back-translation as a domain adaptation problem, simplifying the works of [16, 6, 52] that explicitly differentiate between the two data using noise/tags/gates. Instead of tagging the synthetic data, our approach – the *tag-less back-translation* – aims to enable the model to learn efficiently from the two data through pre-training and fine-tuning. Instead of relying on the model to differentiate between the data, we used the synthetic data as generic domain (out-of-domain) and pre-train the model on this data. We then used the authentic data as in-domain to fine-tune the pre-trained translation model. We hypothesize that although the tagging and noising approaches improve the forward models, our domain-adaptation-tailored approach will provide a flexible method of maximizing the gains in the quantity of the synthetic data and efficiently utilizing the quality in the authentic

parallel data. The approach will also enable the use of different training settings on the different data, as obtainable in domain adaptation strategies. It also gets better as more research and more improved ways of domain adaption are proposed.

In domain adaptation, the generic model is not always expected to perform very well in the domain it is to be deployed, hence the model is fine-tuned with a usually smaller but in-domain data. In many languages, the in-domain data is usually low-resourced or non-existent: having the same issue as in low resource neural machine translation. The in-domain data in itself is not sufficient to create a good model while the more abundant out-of-domain data performs poorly when deployed in the target domain. Mixing the two data results in the in-domain data to be lost in the out-of-domain data and the resulting model is not able to also perform well in the target domain. The larger out-of-domain data is, therefore, used to pre-train a model and the weights of this model are used to initialize the training of the in-domain translation model – a technique referred to as *fine-tuning* [11]. When a different language pair is used for pre-training than that used during fine-tuning, the approach is regarded to as transfer learning [55, 28].

We make the following contributions in this paper:

- we proposed a novel approach that enables a translation model that is trained on synthetic and authentic parallel data to be able to efficiently learn from the the two data, utilizing the different advantages presented by each.
- we successfully applied pre-training and fine-tuning to enable the forward model in back-translation to differentiate between synthetic and authentic data during training, achieving a superior performance to standard and the successful tagged back-translation approaches,
- experimental results have shown that the approach is superior to the standard and tagging back-translation approaches in low resource English-Vietnamese and English-German neural machine translation systems.

The remaining sections are as follows: Section 2 reviews relevant literature on NMT, leveraging monolingual data in NMT and pre-training and fine-tuning. Section 3 explains the tag-less back-translation approach, Section 4 describes the data and experimental set-up used in training the models, Section 5 discusses the results obtained after the experiment. We discuss further the findings in Section 6. Finally, in Section 7, we concluded the work and suggest future directions.

2 Related Works

This section presents prior work on NMT, back-translation and pre-training in NMT.

2.1 Neural Machine Translation (NMT)

The NMT is based on a sequence-to-sequence encoder-decoder system with attention mechanism [2, 47, 34]. The encoders and decoders are made of neural networks that model the conditional probability of a target sentence y given the source sentence x : $p(y|x)$. The encoder converts the input in the source language into a set of vectors while the decoder converts the set of vectors into the target language through an attention mechanism, one word at a time. The attention mechanism was introduced to keep track of context in longer sentences [2].

The NMT model produces the translation sentence by generating one target word at every time step. Given an input sequence $X = (x_1, \dots, x_{T_x})$ and previously translated words $(y_1, \dots, y_{i-1})$, the probability of the next word y_i is

$$p(y_i|y_1, \dots, y_{i-1}, X) = g(y_{i-1}, s_i, c_i) \quad (1)$$

where s_i is the decoder hidden state for time step i and is computed as

$$s_i = f(s_{i-1}, y_{i-1}, c_i). \quad (2)$$

Here, f and g are nonlinear transform functions, which can be implemented as long short-term memory (LSTM) network [23] or gated recurrent units (GRU) [9] in recurrent neural machine translation (RNMT), and c_i is a distinct context vector at time step i , which is calculated as a weighted sum of the input annotations h_j

$$\sum_{j=1}^{T_x} a_{i,j} h_j \quad (3)$$

where h_j is the annotation of x_j calculated by a bidirectional Recurrent Neural Network. The weight $a_{i,j}$ for h_j is calculated as

$$a_{i,j} = \frac{\exp e_{i,j}}{\sum_{t=1}^{T_x} \exp e_{i,t}} \quad (4)$$

and

$$e_{i,j} = v_a \tanh(W s_{i-1} + U h_j) \quad (5)$$

where v_a is the weight vector, W and U are the weight matrices.

All of the parameters in the NMT model, represented as θ , are optimized to maximize the following conditional log-likelihood of the M sentence aligned bilingual samples

$$L(\theta) = \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^{T_y} \log(p(y_i^m | y_{<i}^m, X^m, \theta)) \quad (6)$$

To remove the recurrence and enable parallelization across multiple GPUs during training, the convolutional neural networks were used to create the convolutional NMT (CNMT) encoder-decoder architecture [18, 51]. The CNMT

utilizes 1-dimensional convolutional layers followed by gated linear units, GLU [14]. The decoders compute and apply attention to each of the layers. The model uses positional embeddings along with residual connections [18].

The transformer [48, 15] architecture was introduced to remove the recurrence and convolutions of previous architectures. The transformer is based solely on multi-headed self-attention layers. It enables parallelization across multiple GPUs, thereby, reducing training time. The architecture is used in current state-of-the-art translation systems [16, 38].

In this work, we used a unidirectional LSTM encoder-decoder architecture with Luong attention [34]. This is a simple recurrent neural network RNMT architecture. Our approach is not architecture-dependent and can be applied to the other architectures or other more enhanced implementations of the RNMT.

2.2 Leveraging Monolingual Data for NMT

The use of monolingual data of the target and/or source language has been studied extensively to improve the performance of neural translation models, especially in low resource settings. [20] explored integrating language models trained on monolingual data into NMT systems, [12, 5] proposed augmenting a copy or slightly modified copy respectively of the target data as source, [45] proposed the back-translation approach, [53] proposed the forward translation and [21] used both forward and back-translations to improve the translation models. The back-translation approach has been shown to outperform other approaches in low and high resource languages [16, 22].

Various studies have investigated back-translation to improve the backward model, to select the most suitable generation/decoding methods for generating the synthetic data and to reduce the impact of higher ratio of the synthetic to the authentic bitext. The quality of the models trained using back-translation depends on the quality of the backward model [16, 17, 22, 5, 19, 29, 52]. To improve the quality of the synthetic parallel data, [22] used iterative back-translation – iteratively using the back-translated data to improve both the backward and forward models. [29] and [13] used high resource languages through transfer learning and [54] explored the use of both target and source monolingual data to improve both the backward and forward models. [37] trained a bilingual system based on [25] to do both forward and backward translations, eliminating the need for two separate models. [41] used Transductive data selection methods to select monolingual data that are in the same domain as the test set for back-translation, improving performance.

The works of [17, 42] have found that the ratio of synthetic to authentic data affects the performance of the models most. When the ratio is high, the model tends to learn more from the synthetic data, which contains more mistakes than the authentic data. Investigations have found that the sampling approach of synthetic data generation and adding noise to beam search output outperforms the regular beam decoding technique [16, 24]. These approaches

were said to improve the models by enhancing source-side diversity. [6] claimed, instead, that the noise only indicates to the model that the input is either synthetic or authentic, enabling the model to better utilize the two data. [52] and [6] used tags (and gates) to enable the model to distinguish between the data and the approach has been shown to efficiently utilize more synthetic data, outperforming standard back-translation and enhancing the efficiency of iterative back-translation.

2.3 Domain Adaptation

Domain adaptation is the use of a usually few amount of in-domain data to improve the performance of an out-of-domain (general purpose) model before deployment. The amount of the in-domain parallel data is usually not sufficient to train a very good model and the general purpose models usually performs poorly [32]. There are two categories of domain adaptation – *data centric* and *model centric* [11] with each having several techniques. The techniques in these classifications include using monolingual data [20], synthetic data generation [45], using data selection [49] and using tagged out-of-domain parallel data [10] and fine-tuning [45]

Pre-training has been used successfully in various machine learning tasks to improve performance when the data is not enough to train a good enough model. It was used for training word embeddings [35], in computer vision [50], fine-tuning NMT models [16] and as transfer learning in low resource NMT [55, 28]. The transfer learning for machine translation approach involves training a model on a high resource language pair and transferring the training on a low resource pair. The works of [55, 36, 28] have shown tremendous improvements over models that are trained with the low resource data from scratch.

In back-translation, [45] showed that fine-tuning a pre-trained model on in-domain data improves the quality of back-translated model. [43] pre-trained the model on the authentic data and fine-tunes it on the mixed synthetic and authentic data. [29] and [13] pre-trained a model on a high resource language and fine-tunes it on a low resource language pair.

3 The Proposed Method

The approach is shown in Fig. 1. As illustrated in Algorithm 1, the authentic parallel data: $D^P = \{(x^{(u)}, y^{(u)})\}_{u=1}^U$ is used to train a target-to-source model, $M_{x \leftarrow y}$. This model – the *backward* model – is used to translate the monolingual target data, $Y = \{(y^{(v)})\}_{v=1}^V$, to generate the synthetic parallel data: $D' = \{(x^{(v)}, y^{(v)})\}_{v=1}^V$. Instead of mixing the two data to train a forward (target) model, we used only the synthetic data to pre-train the forward model, $M_{x \rightarrow y}$, until no improvement is observed on the development set. Finally, the forward model is fine-tuned on authentic data.

It was shown in [29] that using different vocabulary each during pre-training and fine-tuning leads to drop in performance because, it was said,

Fig. 1: Tag-less Back-Translation: Training the forward model on the synthetic parallel data generated using the backward model. The forward model is then fine-tuned on the authentic data.

independent vocabulary use different identifiers even for the same subwords and the network loses benefits of the weights learned during pre-training. The authors proposed learning a joint BPE on a mixture of both the pre-training and fine-tuning data and this has been shown to achieve better results in domain adaptation. In this approach, we have access to both the out-of-domain (synthetic) and the in-domain (authentic) parallel data. This, therefore, enables us to learn a joint BPE and build the training vocabulary for both pre-training and fine-tuning.

Algorithm 1: Tag-less Back-Translation

Input: Parallel data $D^P = \{(x^{(u)}, y^{(u)})\}_{u=1}^U$ and
Monolingual target data $Y = \{(y^{(v)})\}_{v=1}^V$

- 1: **procedure** BACK-TRANSLATION
- 2: Train backward model $M_{x \leftarrow y}$ on bilingual data D^P
- 3: Use $M_{x \leftarrow y}$ to create $D' = \{(x^{(v)}, y^{(v)})\}_{v=1}^V$, for $y \in Y$;
- 4: Pre-train forward model $M_{x \rightarrow y}$ on parallel data D' ;
- 5: Fine-tune the forward model $M_{x \rightarrow y}$ on parallel data D^P ;
- 6: **end procedure**

Output: forward model $M_{x \rightarrow y}$

Table 1: Data Used

data	train		dev	test
	sentences	words (vocab)		
		EnVi		
IWSLT'15 En-Vi	133, 317	2,837,240 (50,045)2,688,387 (103,796)	1, 553	1, 268
IWSLT'14 En-De	153, 348	EnDe2,706,255 (54,169)3,311,508 (25,615)	6, 970	6, 750
WMT'14 En-De – Monolingual English	666,585	16,700,106 (344,138)	-	-

4 Experiments

4.1 Set-up

We used the TensorFlow [1] implementation of the OpenNMT [27] framework to train the models – the NMTSmallV1 configuration. The set-up is based on the NMTSmallV1 configuration. Specifically, the configuration is a 2-layer unidirectional LSTM encoder-decoder model with Luong attention [34]. It has 512 hidden units and a vocabulary size of 50,000 for both source and target languages. We used Adam [26] optimizer and a batch size of 64 with a dropout probability of 0.3, a static learning rate of 0.0002 and the models are evaluated on the development set after every 5,000 training steps. The models were evaluated using the bi-lingual evaluation understudy metric, BLEU [39], specifically the multi-bleu [31] implementation. The models are trained until there is no improvement of over 0.2 BLEU after four training steps. As stated in Section 3, the learning of BPE on the training data and the building of training vocabulary for both pre-training and fine-tuning was done on the mixture of the synthetic and authentic parallel data. During fine-tuning, we, therefore, only change the training data.

4.2 Data

For this work, we use the preprocessed low resource English-Vietnamese parallel data [34] of the IWSLT 2015 Translation task [7]. We used the 2012 and 2013 test sets for development and testing respectively. We also used the data from the IWSLT 2014 German-English shared translation task [8] as the second language pair, pre-processed using the data clean-up as well as the train, development and test split in [44]. For the monolingual data, we used the pre-processed English monolingual data of WMT 2014 English-German translation task [4]. We shuffled the monolingual data and selected 666,585 monolingual sentences which is five times as much as the En-Vi parallel data. The statistics

of the datasets are shown in Table 1. We learned byte pair encoding (BPE) [46] with 10,000 merge operations on the training dataset and applied it on the train, development and test datasets. Afterwards, we build the vocabulary on the training dataset. For all the experiments, we used thrice as much of the monolingual as the available parallel data in both of the languages except when we experimented with the ratio of 1:5 (parallel to monolingual data) for the English-Vietnamese NMT.

4.3 Models

To compare the performance of our approach with that of the previous works, we implemented the following methods to train translation models on the English-Vietnamese and English-German NMT using the data presented in Section 4.2 above. All models were trained using the same settings stated in Section 4.1

- We first train *baseline* models using the available authentic parallel data only. In the models, the baselines have the English language as the target language – Vi-En and De-En.
- We then train the backward models also on the authentic parallel data using English language as the source language – En-Vi and En-De. The models are used for the generation of the additional synthetic parallel data for the back-translation approach.
- We implemented the various back-translation strategies namely standard back-translation – *standard_bt*, the tagged back-translation – *tagged_bt* and the tag-less back-translation – *tag-less_bt* (*joint BPE*) using the authentic and synthetic parallel data.

5 Results

All scores reported are statistically significant with $p < 0.05$. We used the paired bootstrap resampling of [30] as implemented in [40] to estimate the statistical significance confidence scores. See Table 13 in Appendix 1 for confidence scores.

5.1 English-Vietnamese Low Resource NMT

The evaluation scores of the best models and the improved models obtained after taking the checkpoint averaging of the last 8 checkpoints are shown in Table 2. We first created a forward Vietnamese-English (Vi-En) model, *baseline*, on the available authentic parallel data. The model trained for 75,000 steps before the stopping condition was met. The baseline model was trained further to 110,000 training steps but the performance continued to flatten without observing any improvement. The model achieved the best single-checkpoint

Table 2: Performance of the Tag-less Back-translated model compared to the baseline and standard back-translation models for Vietnamese-English and German-English translations. Evaluation scores on the test set. The tag-less approaches show results of pre-training and fine-tuning.

data		<i>baseline</i>	<i>standard_bt</i>	<i>tag-less_bt</i>	<i>tag-less_bt</i> (joint BPE)
Vi-En	best checkpoint				
	BLEU (training step)	21.19 (65k)	24.46 (155k)	17.85 (90k) 25.16 (145k)	18.60 (105k) 26.53 (165k)
	average	22.22	25.28	17.96 25.77	18.59 26.83
De-En	best checkpoint				
	BLEU (training step)	20.30 (75k)	25.11 (150k)	3.01 (75k) 25.16 (120k)	5.43 (60k) 28.31 (155k)
	average	20.95	25.87	3.03 26.03	5.13 28.83

score of 21.19 BLEU at the 65,000th. We then trained a backward (En-Vi) model, *backward*. After the stoppage condition were met, after 55,000 training steps, the best performing single-checkpoint for the backward model achieved a BLEU score of 24.78 at the 50,000th training step. Averaging the last 8 checkpoints gave the best performance – 25.79 BLEU. This average model was used to back-translate the monolingual English data to generate synthetic parallel data. We mixed the two data – synthetic and authentic – without differentiating between the two and used the resulting large dataset to train a forward model. We labelled this model as *standard_bt* – for standard back-translation [45]. This model was trained for 165,000 before the stopping condition were met. It achieved a single-checkpoint best BLEU score of 24.46 at the 155,000th training step.

We mixed the synthetic and authentic parallel data and learned a joint BPE on the resulting training dataset and build the training vocabulary. We applied the BPE on the synthetic data for pre-training and on the authentic data for fine-tuning. We trained a model, labelled *tag-less_bt (joint BPE)*, using this approach. The model achieved a single-checkpoint score of 18.60 BLEU during pre-training and improved to 26.53 BLEU after fine-tuning. The average fine-tuned model was better by about 0.30 BLEU. The average pre-trained model performed very low compared to the baseline and the standard back-translation models – 18.59 BLEU vs 22.22 and 25.28 BLEUs respectively. This is obviously because the quality of the data used in the training the model – the synthetic data – is lower than that of the other two. The quality of the synthetic data, although generated from a reasonably good backward model, is still not sufficient to train a model whose quality can compare to the other models that are trained in whole or in part with the authentic data. Fine-tuning the model on the authentic data results in a sharp rise in performance.

Fig. 2: Tag-less back-translation: pre-training on synthetic data and fine-tuning on authentic data. Showing how this technique compares to the baseline and the standard back-translation approaches on the test set.

The model was fine-tuned until the stopping condition was met. The approach outperformed the baseline and standard back translation models by 5.34 and 2.07 BLEUs respectively. The gap in performance was, however, reduced to 4.61 and 1.55 BLEUs after checkpoint averaging.

We experimented the other pre-train and fine-tune approach, learning the BPE only on the synthetic data. We build the vocabulary on the synthetic training data after applying the BPE. The synthetic corpus was used to pre-train a forward model for 130,000 steps, achieving a single-checkpoint best score of 17.85 BLEU. The authentic parallel data was then used to fine-tune the model for a further 35,000 training steps. Stopping at each of these steps were based on the stopping condition. The performance of the *tag-less_bt* model improved to 25.16 BLEU after fine-tuning. Although this approach was shown to outperform the baseline and standard back-translation, it underperformed the joint BPE implementation of the tag-less approach by 1.06 BLEU. In Figures 2a and 2b, we show how the BLEU scores continue to improve with increase in the training steps. The model trained using the *tag-less_bt (joint BPE)* approach continued to outperform the three others after fine-tuning.

5.2 English-German Low Resource NMT

We conducted the same set of experiments presented in section 5.1 on the second low resource dataset, the English-German IWSLT’14 parallel dataset. This data, as presented in Table 1, is made up of a little bit more than 150,000 parallel sentences. We first trained a backward (En-De) model on the available parallel data. This model maxed-out performance on the test set, based on the set-up, at 10.25 BLEU after averaging the last 8 checkpoints. It stopped training at the 80,000th training steps and achieving the best single model performance at the 65,000th – 10.03 BLEU. We used the average model to generate the synthetic data, translating the available English monolingual data. We

Fig. 3: Tagged Vs Tag-less back-translation.

trained four separate forward (De-En) models based on the approaches we explained earlier. The first is the baseline trained on the available authentic data, the *standard_bt* on the mixture of the authentic and synthetic data without differentiating, the *tag-less_bt* pre-trained on the synthetic data and fine-tuned on the authentic data having learned the BPE on the synthetic data and updating the vocabulary before fine-tuning and, finally, the *tag-less_bt (joint BPE)* trained also using the tag-less approach but having learned the BPE and built the vocabulary on the mixture of the synthetic and authentic data.

The results of evaluating the models after training using the various approaches are presented in Table 2. The baseline achieved a modest average performance of 20.95 BLEU after training for 100,000 training steps. The performance on the dataset was improved after applying standard back-translation, achieving a huge +4.92 BLEU improvement over the baseline. The tag-less approach, though better, did not achieve a huge improvement over back-translation (only +0.16 BLEU) but after applying the improved tag-less (joint BPE), as shown in the previous section, we achieved huge +2.96 BLEU increase in performance. This +2.8 and +7.88 BLEUs over the previous tag-less approach and the baseline respectively.

For all the subsequent experiments, unless stated otherwise, we used the joint BPE technique to implement the tag-less back-translation approach as it is shown to be the most successful variant.

5.3 Tagged Vs Tag-Less Back-translation

We compared the performance of the *tag-less_bt* model – our technique – with that of the successful tagged back-translation of [6] on the English-Vietnamese data. The synthetic sources were labelled with the $\langle BT \rangle$ token at the beginning of each sentence and mixed with the authentic sources to generate the mixed tagged parallel corpus. This mixed data is used to train the forward tagged back-translation model – *tagged_bt*. The *tagged_bt* model stopped at

Table 3: Performance of Tag-less Back-translation compared to the Tagged back-translation model for Vietnamese-English translation. Evaluation scores on the test set.

		<i>tagged_bt</i>	<i>tag-less_bt (joint BPE)</i>
En-Vi	best checkpoint BLEU (training step)	24.76 (165k)	26.53 (165k)
	average	25.05	26.83
En-De	best checkpoint BLEU (training step)	27.49 (205k)	28.31 (155k)
	average	27.75	28.83

Table 4: Performance of Tag-less Back-translation on the test set: pre-training on synthetic data and fine-tuning on authentic data Vs pre-training on authentic data and fine-tuning on synthetic data for Vietnamese-English.

	<i>tag-less_bt</i>	<i>reverse_tag-less_bt</i>
best checkpoint BLEU (training step)	25.16 (145k)	21.19 (65k) - pre-train 18.95 (100k) - fine-tune
average	25.77	18.91

125,000 steps and the training was continued up to 195,000 steps to equal the number of training steps reached by the *tag-less_bt* model. While the tagged approach underperformed the best score of our technique by 1.78 BLEU, it was able to outperform the single-checkpoint standard back-translation by 24.76 to 24.46 BLEUs respectively (+0.3 BLEU) but underperformed the average standard back-translation model by 0.23 BLEU.

Finally, we trained a forward model using the tagged back-translation for English-German NMT to compare the performance with our approach on this data. The tagged approach took a further 50,000 training steps to reach a single model best of 27.49 BLEU, but still underperforming the tag-less approach by 0.82 BLEU. The best model obtained after averaging checkpoints was also achieved using our approach, a performance of 28.83 BLEU compared to the tagged 27.75 BLEU, an improvement of 1.08 BLEU. The performances of these models, evaluated on the test set is shown in Table 3. On this data, the tagged approach performed better than the standard back-translation by +1.88 BLEU on the average models. It can be seen in Figures 2 and 3 that in both of the experiments conducted on the two data, our tag-less approach out-performed the rest of the back-translation approaches.

This supports the hypothesis that although the tagged back-translation involves explicit differentiating between the two data using tags, the model trained on the approach may not be able to differentiate between them com-

Fig. 4: Fine-tuning the baseline model on the synthetic data. Evaluation scores on the test set.

pletely during training as observed in the mixed performance of the models trained on the two different data.

5.4 Fine-tuning: Synthetic Vs Authentic Data

Our technique proposed pre-training the forward model on the synthetic parallel data and fine-tuning the model afterwards on the authentic data. This was proposed to enable the model to unlearn the mistakes it learned from the synthetic data using correct sentences in the authentic parallel data. We experiment the other way round to investigate the effects of pre-training on the authentic data and fine-tuning on the synthetic data. We used the baseline as the pre-trained model and fine-tune it on the synthetic data. This approach was labelled as *reverse_tag-less_bt*. This approach did not show any benefit to the final forward model, see Fig. 3. As expected, the performance of the model decreased and the curve flattens as the number of training steps increases. The best and average scores are shown in Table 3.

5.5 Quantity of Monolingual Data

As stated earlier, it was found that as the more synthetic data increases, a point is reached where the performance starts to deteriorate [17]. Instead, our work hypothesizes that the performance of the model will start to decrease only if it is not able to differentiate between the synthetic and authentic training data and, therefore, efficiently learning from the two. We also pointed out

Table 5: Using different ratios of the authentic to synthetic data for Vietnamese-English translation. Evaluation scores of the models on the test set.

	<i>tagged.bt</i>			<i>tag-less.bt</i> (joint BPE)		
	1:1	1:3	1:5	1:1	1:3	1:5
best checkpoint BLEU (training step)	22.73 (85k)	24.76 (165k)	24.62 (155k)	26.29 (130k)	26.53 (165k)	26.91 (185k)
average	22.73	25.05	25.47	26.15	26.83	27.14

that since the data is mixed in both the standard and tagged back-translation approaches, the model may not be able to completely differentiate between the data, although in the latter approach, the model is expected to treat the tagged synthetic sources as a different domain. We, therefore, experiment with different ratios of the authentic to synthetic data to verify this claim. We sample the authentic to synthetic data in the ratios, 1:1, 1:3 and 1:5. The results are shown in Table 5.

In the tagged approach, the single-checkpoint best scores continue to rise from using the same amount of monolingual data for back-translation to using three times the authentic data of the monolingual data for back-translation. But, as observed, the performance dropped slightly when we used five times the amount of available parallel data. However, the performance of the tag-less back-translation models continues to increase steadily when the ratio of authentic to monolingual data is increased. We observed the performance to improve by about 0.25 BLEU when the amount of synthetic data is tripled and doubled to about 0.5 BLEU after adding another double amount of the synthetic data to the training data. It can also be observed that there was a very low improvement over the performance of *baseline* and serious underperformance compared to the tag-less approach when we used the same amount of synthetic data with the authentic data to train the models – 22.73 BLEU vs 22.22 and 26.15 BLEUs respectively. Overall, we obtained a 3.42, 1.78 and 1.67 BLEU improvements on the average models using the tag-less approach over the tagged approach on the ratios experimented respectively.

It can be observed also that the performance of the model trained on the 1:1 ratio of monolingual to synthetic data using our approach is very good compared to the model trained using the same amount of data in the tagged approach and subsequent increase in training data leads to steady improvements that at 1:5 ratio, the performance was improved by about 1 BLEU. This steady improvements can show that the model learned useful knowledge on the authentic data but only used the synthetic data for further improvements. Following this trend, we can, therefore, conclude that with more synthetic data compared to the authentic data, the model will only continue

Table 6: Before and after fine-tuning the English-Vietnamese standard and tagged back-translation NMT models on the authentic data. Evaluation scores on the test set.

	<i>tagged_bt (1:5)</i>		<i>standard_bt (1:3)</i>	
	before	after	before	after
best checkpoint BLEU (training step)	24.62 (155k)	25.55 (175k)	24.46 (155k)	25.15 (180k)
average	25.47	25.64	25.28	25.32

to learn and increase its performance if useful representations are learnable on the synthetic data.

5.6 Fine-tuning Standard And Tagged Back-Translations

The work of [43] reported no observable advantage of using the authentic data to train the forward model and fine-tuning it henceforth on the mixed data. Instead, we experimented training the forward model on the mixed data first and then fine-tune it on the authentic data. As shown in Table 6, this approach reaches the same performance as the old tag-less approach – 25.15 BLEU – using the same amount of synthetic sentences albeit after 30,000 more training steps but still underperforming the joint BPE tag-less back-translation’s 26.53 BLEU although training for additional 15,000 training steps. The better joint BPE tag-less approach converges earlier than fine-tuning the standard back-translation model, at 165,000 training steps.

We also explored the use of fine-tuning to determine whether or not the tagged approach will be able to cover the difference in performance with the tag-less approach. After fine-tuning the *tagged_bt (1:5)* model for a further 35,000 training steps, the performance gained was a significant +0.93 BLEU and only 0.17 BLEU over the average after just 20,000 steps of fine-tuning. The performance was still short of the *tag-less_bt (joint BPE) (1:5)* by a significant 1.36 BLEU.

6 Discussion

In this work, we proposed an approach for training the forward model in back-translation without using tags or noising the synthetic data. Translation models that are trained on the synthetic and authentic data have been shown to perform better when they are able to differentiate between the two data. Previous approaches have relied on the use of noise in back-translation [16] especially on low resource languages to improve the performance of models. The authors thought that the approach ensures source-side diversity which has been shown to benefit the models [24]. The approach was found out to

only indicate to the forward model that the noised data is synthetic, enabling it to treat the data differently from the authentic data [6]. The use of tags has been shown to improve the performance of such models. In this work, we eliminated the need of using of the tags and showed that although it was successful at improving the performance – proving it successful at indicating to the model that a data is synthetic and not authentic – domain adaptation methods are more capable of ensuring the model differentiate between the data. The ability for the model to separate between the data is even more important in low resource languages where the available data is not enough to train standard backward model, thus generating synthetic data with a lot of noises.

Domain adaptation techniques techniques in machine translation ensures that a better model is trained, leveraging on a larger parallel data of either the same language pair but in a different domain – fine-tuning – or a different language pair – transfer learning. In this technique, the two data are not tagged, mixed and left to the model to differentiate between them. They, rather, are used at the different stages of the training and this ensures the model performs in the target domain as expected. We utilized the synthetic data – which is bigger but more prone to translation noises – as the generic domain and the authentic data – smaller but having more quality – as the in-domain. This selection was not done until the reversed approach was shown not give the desired performance. The superiority of the approach over the successful tagging was shown through experimental results conducted on two low resource language pair: English-Vietnamese and English-German. In each of the languages considered, we obtained an improvement of more than 1 BLEU points over the tagged approach that outperformed the baseline and standard back-translation models.

We also test the performance of our technique when the amount of monolingual data is increased. We used different ratios of the authentic parallel to monolingual data used. We found that our technique was not only able to handle the increase in the synthetic data, but was able to attain rapid improvement given the smallest amount of synthetic data. We obtained a superior performance by a whopping 3.56 BLEU using the tag-less approach over the tagged approach when the amount of monolingual data is the same as the authentic data. The performance continued to steady increase as the amount of monolingual data is increased. The tagged approach could only handle tripling the amount of synthetic data but the performance started to decrease when the synthetic data was increased further. Using the same amount of synthetic data in ratios 1:1, 1:3 and 1:5, our technique outperformed the tagging technique by 3.42, 1.78 and 1.67 respectively (see average scores in Table 5).

Our approach also provides one with the flexibility of using state-of-the-art domain adaptation methods to improve the performance of the already successful back-translation approach. Techniques such as using different dropout and/or learning rate during pre-training and fine-tuning may improve the performance of the forward model. The method may also be applied in high-resource languages since both of these settings – low and high resource – can

benefit from the ability of the forward model to differentiate between synthetic and authentic data.

7 Conclusions and Future Work

This work has shown that an NMT model pre-trained on synthetic data and fine-tuned on the authentic data outperforms the rather successful method of tagging the synthetic data in low resource NMT by enabling the forward model to differentiate between the authentic and synthetic training data. The approach, however, does not improve the performance when it is reversed and the forward model is pre-trained on the authentic data and then fine-tuned on the synthetic data. As expected, the reverse approach makes the model to unlearn the useful representations learned in favour of the noise in the synthetic data. This justifies our hypothesis that without differentiating between the two data, the synthetic data is most likely to hurt the performance of the forward model.

It was shown also, in this work, that the more synthetic data used, the better the performance of the forward model, though the most effective ratio was not yet determined through thorough experimentation. This will inform the basis of future works. We experiment fine-tuning the models trained using the standard and tagged back-translation approaches. Experimental results showed the standard back-translation equalling the performance of a variant of the tag-less approach after many more rounds of training. The performance of the tagged approach improved considerably but still trailed the tag-less approach. The most successful of the tag-less approach has been the one that involves learning a joint BPE and building the training vocabulary on the mixture of the synthetic and authentic parallel data. This approach is made possible, unlike in other fine-tuning conditions, because both the generic (synthetic) data and the in-domain (authentic) data are available during the process.

For future work, the use of different settings – such as increasing or decreasing the learning rate, using dropout and L2 regularization, which may reduce overfitting on the in-domain (authentic) data as shown to be a likely problem in domain adaptation by [3] – for the pre-training and fine-tuning approaches can be explored to maximize the benefits of the domain adaptation approach in back-translation. The approach can also be investigated to improve the forward translation approach – which also leverages on the synthetic data for additional training data. Finally, we intend to investigate the technique in high resource languages in the future.

Conflict of interest

The authors declare that they have no conflict of interest.

References

1. Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., Kudlur, M., Levenberg, J., Monga, R., Moore, S., Murray, D.G., Steiner, B., Tucker, P., Vasudevan, V., Warden, P., Wicke, M., Yu, Y., Zheng, X.: Tensorflow: A system for large-scale machine learning. In: Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation, OSDI'16, p. 265–283. USENIX Association, USA (2016)
2. Bahdanau, D., Cho, K., Bengio, Y.: Neural Machine Translation by Jointly Learning to Align and Translate. In: Y. Bengio, Y. LeCun (eds.) 3rd International Conference on Learning Representations, {ICLR} 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015). URL <http://arxiv.org/abs/1409.0473>
3. Barone, A.V.M., Haddow, B., Hermann, U., Sennrich, R.: Regularization techniques for fine-tuning in neural machine translation. In: Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, pp. 1489–1494. Association for Computational Linguistics, Copenhagen, Denmark (2017). DOI 10.18653/v1/D17-1156. URL <https://www.aclweb.org/anthology/D17-1156>
4. Bojar, O., Buck, C., Federmann, C., Haddow, B., Koehn, P., Leveling, J., Monz, C., Pecina, P., Post, M., Saint-Amand, H., Soricut, R., Specia, L., Tamchyna, A.: Findings of the 2014 Workshop on Statistical Machine Translation. In: Proceedings of the Ninth Workshop on Statistical Machine Translation, pp. 12–58. Association for Computational Linguistics, Baltimore, Maryland, USA (2014). DOI 10.3115/v1/w14-3302
5. Burlot, F., Yvon, F.: Using Monolingual Data in Neural Machine Translation: a Systematic Study. In: Proceedings of the Third Conference on Machine Translation: Research Papers, pp. 144–155. Association for Computational Linguistics, Brussels, Belgium (2018). DOI 10.18653/v1/W18-6315. URL <https://www.aclweb.org/anthology/W18-6315>
6. Caswell, I., Chelba, C., Grangier, D.: Tagged Back-Translation. In: Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers), pp. 53–63. Association for Computational Linguistics, Florence, Italy (2019). DOI 10.18653/v1/W19-5206. URL <https://www.aclweb.org/anthology/W19-5206>
7. Cettolo, M., Federico, M., Bentivogli, L., Niehues, J., Stüker, S., Sudoh, K., Yoshino, K., Federmann, C.: Overview of the IWSLT 2017 Evaluation Campaign. In: Proceedings of the 14th International Workshop on Spoken Language Translation. Tokyo, Japan (2017)
8. Cettolo, M., Niehues, J., Stüker, S., Bentivogli, L., Federico, M.: Report on the 11th IWSLT Evaluation Campaign, IWSLT 2014. In: Proceedings of the 11th Workshop on Spoken Language Translation, pp. 2–16. Lake Tahoe, CA, USA (2014)
9. Cho, K., van Merriënboer, B., Gülçehre, Ç., Bougares, F., Schwenk, H., Bengio, Y.: Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. CoRR **abs/1406.1** (2014). URL <http://arxiv.org/abs/1406.1078>
10. Chu, C., Dabre, R., Kurohashi, S.: An empirical comparison of domain adaptation methods for neural machine translation. In: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), pp. 385–391. Association for Computational Linguistics, Vancouver, Canada (2017). DOI 10.18653/v1/P17-2061. URL <https://www.aclweb.org/anthology/P17-2061>
11. Chu, C., Wang, R.: A survey of domain adaptation for neural machine translation. In: Proceedings of the 27th International Conference on Computational Linguistics, pp. 1304–1319. Association for Computational Linguistics, Santa Fe, New Mexico, USA (2018). URL <https://www.aclweb.org/anthology/C18-1111>
12. Currey, A., Miceli Barone, A.V., Heafield, K.: Copied Monolingual Data Improves Low-Resource Neural Machine Translation. In: Proceedings of the Second Conference on Machine Translation, vol. 1, pp. 148–156. Association for Computational Linguistics, Copenhagen, Denmark (2017). DOI 10.18653/v1/W17-4715
13. Dabre, R., Chen, K., Marie, B., Wang, R., Fujita, A., Utiyama, M., Sumita, E.: NICT’s Supervised Neural Machine Translation Systems for the WMT19 News Translation Task. In: Proceedings of the Fourth Conference on Machine Translation (WMT), vol. 2, pp. 168–174. Florence, Italy (2019)

14. Dauphin, Y.N., Fan, A., Auli, M., Grangier, D.: Language Modeling with Gated Convolutional Networks. In: Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17, pp. 933–941. JMLR.org (2017)
15. Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., Kaiser, L.: Universal Transformers. ICLR pp. 1–23 (2019). URL <http://arxiv.org/abs/1807.03819>
16. Edunov, S., Ott, M., Auli, M., Grangier, D.: Understanding Back-Translation at Scale. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 489–500. Association for Computational Linguistics, Brussels, Belgium (2018). DOI 10.18653/v1/D18-1045. URL <https://www.aclweb.org/anthology/D18-1045>
17. Fadaee, M., Monz, C.: Back-Translation Sampling by Targeting Difficult Words in Neural Machine Translation. In: Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, pp. 436–446. Association for Computational Linguistics, Brussels, Belgium (2018). DOI 10.18653/v1/D18-1040. URL <https://www.aclweb.org/anthology/D18-1040>
18. Gehring, J., Michael, A., Grangier, D., Yarats, D., Dauphin, Y.N.: Convolutional Sequence to Sequence Learning. In: D. Precup, Y.W. Teh (eds.) Proceedings of the 34th International Conference on Machine Learning, pp. 1243–1252. Sydney, Australia (2017). URL <http://proceedings.mlr.press/v70/gehring17a.html>
19. Graça, M., Kim, Y., Schamper, J., Khadivi, S., Ney, H.: Generalizing Back-Translation in Neural Machine Translation. In: Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers), pp. 45–52. Association for Computational Linguistics, Florence, Italy (2019). DOI 10.18653/v1/W19-5205. URL <https://www.aclweb.org/anthology/W19-5205>
20. Gulcehre, C., Firat, O., Xu, K., Cho, K., Bengio, Y.: On integrating a language model into neural machine translation. *Computer Speech & Language* **45**(2017), 137–148 (2017). DOI 10.1016/j.csl.2017.01.014
21. He, D., Xia, Y., Qin, T., Wang, L., Yu, N., Liu, T.Y., Ma, W.Y.: Dual Learning for Machine Translation. In: Proceedings of the 30th International Conference on Neural Information Processing Systems, NIPS'16, pp. 820–828. Curran Associates Inc., USA (2016). URL <http://dl.acm.org/citation.cfm?id=3157096.3157188>
22. Hoang, V.C.D., Koehn, P., Haffari, G., Cohn, T.: Iterative Back-Translation for Neural Machine Translation. In: Proceedings of the 2nd Workshop on Neural Machine Translation and Generation, pp. 18–24. Association for Computational Linguistics, Melbourne, Australia (2018)
23. Hochreiter, S., Schmidhuber, J.: Long Short-Term Memory. *Neural Computation* **9**(8), 1735–1780 (1997)
24. Imamura, K., Fujita, A., Sumita, E.: Enhancement of Encoder and Attention Using Target Monolingual Corpora in Neural Machine Translation. In: Proceedings of the 2nd Workshop on Neural Machine Translation and Generation, pp. 55–63. Association for Computational Linguistics, Melbourne, Australia (2018)
25. Johnson, M., Schuster, M., Le, Q.V., Krikun, M., Wu, Y., Chen, Z., Thorat, N., Viégas, F., Wattenberg, M., Corrado, G., Hughes, M., Dean, J.: Google's Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation. *Transactions of the Association for Computational Linguistics* **5**, 339–351 (2017). URL <http://arxiv.org/abs/1611.04558>
26. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimization. In: Y. Bengio, Y. LeCun (eds.) 3rd International Conference on Learning Representations, {ICLR} 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings (2015). URL <http://arxiv.org/abs/1412.6980>
27. Klein, G., Kim, Y., Deng, Y., Senellart, J., Rush, A.M.: OpenNMT: Open-Source Toolkit for Neural Machine Translation. *Proceedings of {ACL} 2017, System Demonstrations* pp. 67–72 (2017). URL <https://www.aclweb.org/anthology/P17-4012>
28. Kocmi, T., Bojar, O.: Trivial Transfer Learning for Low-Resource Neural Machine Translation. In: Proceedings of the Third Conference on Machine Translation (WMT), vol. 1, pp. 244–252. Brussels, Belgium (2018). DOI 10.18653/v1/W18-6325
29. Kocmi, T., Bojar, O.: CUNI Submission for Low-Resource Languages in WMT News 2019. In: Proceedings of the Fourth Conference on Machine Translation (WMT), vol. 2, pp. 234–240. Florence, Italy (2019)

30. Koehn, P.: Statistical significance tests for machine translation evaluation. In: Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing, pp. 388–395. Association for Computational Linguistics, Barcelona, Spain (2004). URL <https://www.aclweb.org/anthology/W04-3250>
31. Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., Herbst, E.: Moses: Open Source Toolkit for Statistical Machine Translation. In: Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions, ACL '07, pp. 177–180. Association for Computational Linguistics, Stroudsburg, PA, USA (2007). URL <http://dl.acm.org/citation.cfm?id=1557769.1557821>
32. Koehn, P., Knowles, R.: Six Challenges for Neural Machine Translation. In: Proceedings of the First Workshop on Neural Machine Translation, pp. 28–39. Association for Computational Linguistics, Vancouver, Canada (2017). DOI 10.18653/v1/w17-3204
33. Koehn, P., Och, J.F., Marcu, D.: Statistical Phrase-Based Translation. In: Proceedings of the 2003 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics, pp. 127–133 (2003). URL <https://www.aclweb.org/anthology/N031017>
34. Luong, M.T., Pham, H., Manning, C.D.: Effective Approaches to Attention-based Neural Machine Translation. In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, pp. 1412–1421. Association for Computational Linguistics, Lisbon, Portugal (2015). DOI 10.18653/v1/D15-1166. URL <https://www.aclweb.org/anthology/D15-1166>
35. Mikolov, T., Corrado, G., Chen, K., Dean, J.: Efficient Estimation of Word Representations in Vector Space. In: Y. Bengio, Y. LeCun (eds.) 1st International Conference on Learning Representations, {ICLR}. Arizona, USA (2013). URL <http://arxiv.org/abs/1301.3781>
36. Nguyen, T.Q., Chiang, D.: Transfer Learning across Low-Resource, Related Languages for Neural Machine Translation. In: Proceedings of the Eighth International Joint Conference on Natural Language Processing, vol. 2, pp. 296–301. Asian Federation of Natural Language Processing (2017)
37. Niu, X., Denkowski, M., Carpuat, M.: Bi-Directional Neural Machine Translation with Synthetic Parallel Data. In: Proceedings of the 2nd Workshop on Neural Machine Translation and Generation, pp. 84–91. Association for Computational Linguistics, Melbourne, Australia (2018). DOI 10.18653/v1/W18-2710. URL <https://www.aclweb.org/anthology/W18-2710>
38. Ott, M., Edunov, S., Grangier, D., Auli, M.: Scaling Neural Machine Translation. In: Proceedings of the Third Conference on Machine Translation: Research Papers, pp. 1–9. Association for Computational Linguistics, Brussels, Belgium (2018). DOI 10.18653/v1/W18-6301. URL <https://www.aclweb.org/anthology/W18-6301>
39. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a Method for Automatic Evaluation of Machine Translation. In: Linguistics, Proceedings of the 40th Annual Meeting of the Association for Computational, pp. 311–318. Association for Computational Linguistics, Philadelphia, Pennsylvania, USA (2002). DOI 10.3115/1073083.1073135. URL <https://www.aclweb.org/anthology/P02-1040>
40. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S.: Pytorch: An imperative style, high-performance deep learning library. In: H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (eds.) Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019). URL <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
41. Poncelas, A., Maillette de Buy, G.W., Way, A.: Adaptation of Machine Translation Models with Back-translated Data using Transductive Data Selection Methods. CoRR **abs/1906.0** (2019). URL <http://arxiv.org/abs/1906.07808>
42. Poncelas, A., Shterionov, D., Way, A., Maillette de Buy, G.W., Passban, P.: Investigating Backtranslation in Neural Machine Translation. CoRR **abs/1804.0** (2018). URL <http://arxiv.org/abs/1804.06189>

43. Popel, M.: Machine Translation Using Syntactic Analysis. Doctoral, Charles University (2018)
44. Ranzato, M., Chopra, S., Auli, M., Zaremba, W.: Sequence Level Training with Recurrent Neural Networks. In: International Conference on Learning Representations (2016)
45. Sennrich, R., Haddow, B., Birch, A.: Improving Neural Machine Translation Models with Monolingual Data. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, pp. 86–96. Association for Computational Linguistics, Berlin, Germany (2016)
46. Sennrich, R., Haddow, B., Birch, A.: Neural Machine Translation of Rare Words with Subword Units. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pp. 1715–1725. Association for Computational Linguistics, Berlin, Germany (2016). URL <https://www.aclweb.org/anthology/P16-1162>
47. Sutskever, I., Vinyals, O., Le, Q.V.: Sequence to Sequence Learning with Neural Networks. In: NIPS (2014)
48. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention Is All You Need. In: 31st Conference on Neural Information Processing Systems. Long Beach, CA, USA (2017)
49. Van der Wees, M., Bisazza, A., Monz, C.: Dynamic data selection for neural machine translation. In: EMNLP 2017 - Conference on Empirical Methods in Natural Language Processing, Proceedings, pp. 1400–1410 (2017). DOI 10.18653/v1/d17-1147
50. Whatmough, P.N., Zhou, C., Hansen, P., Venkataramanaiah, S.K., Seo, J.s., Mattina, M.: FixyNN: Efficient Hardware for Mobile Computer Vision via Transfer Learning. CoRR **abs/1902.1** (2019). URL <http://arxiv.org/abs/1902.11128>
51. Wu, F., Fan, A., Baevski, A., Dauphin, Y.N., Auli, M.: Pay Less Attention with Lightweight and Dynamic Convolutions. In: 7th International Conference on Learning Representations, {ICLR}. New Orleans, LA, USA (2019). URL <https://openreview.net/forum?id=SkVh09tX>
52. Yang, Z., Chen, W., Wang, F., Xu, B.: Effectively training neural machine translation models with monolingual data. *Neurocomputing* **333**, 240–247 (2019). DOI 10.1016/j.neucom.2018.12.032
53. Zhang, J., Zong, C.: Exploiting Source-side Monolingual Data in Neural Machine Translation. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, pp. 1535–1545. Association for Computational Linguistics, Austin, Texas (2016). DOI 10.18653/v1/d16-1160
54. Zhang, Z., Liu, S., Li, M., Zhou, M., Chen, E.: Joint Training for Neural Machine Translation Models with Monolingual Data. CoRR **abs/1803.00353** (2018). URL <http://arxiv.org/abs/1803.00353>
55. Zoph, B., Yuret, D., May, J., Knight, K.: Transfer Learning for Low-Resource Neural Machine Translation. In: Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, pp. 1568–1575. Association for Computational Linguistics, Austin, Texas (2016). DOI 10.18653/v1/D16-1163. URL <https://www.aclweb.org/anthology/D16-1163>

Appendix

In this section, we provided the complete evaluation scores on the test set for all the models trained.

Table 7: Performance of Tag-less Back-translation compared to the baseline and standard back-translation models. BLEU Scores for each Checkpoint of the Models for Vietnamese-English NMT (best single-checkpoint and average scores are shown in **bold**).

training step (thousands)	<i>baseline</i>	<i>standard_bt</i>	<i>tag-less_bt</i>		<i>tag-less_bt (joint BPE)</i>	
			pre-train	fine-tune	pre-train	fine-tune
5	2.50	1.92	2.88		2.87	
10	4.25	3.50	7.54		9.78	
15	5.54	11.34	10.70		13.06	
20	6.68	15.83	12.86		14.11	
25	10.87	17.20	12.96		15.03	
30	16.89	19.47	14.07		15.92	
35	19.47	18.95	15.10		15.95	
40	20.37	20.81	15.59		16.39	
45	21.09	21.61	15.61		17.08	
50	20.13	22.18	15.99		17.15	
55	21.09	21.98	16.69		17.29	
60	21.11	22.35	16.66		17.15	
65	21.19	22.37	16.73		17.37	
70	20.17	23.43	17.63		17.51	
75	19.41	22.92	17.13		17.92	
80	20.43	23.41	17.33		17.96	
85	20.18	23.62	16.97		17.94	
90	19.80	23.71	17.85		18.58	
95	20.36	23.77	17.30		18.28	
100	19.48	23.73	17.31		18.23	
105	19.63	24.07	17.36		18.60	
110	19.77	23.99	17.34		18.53	
115	-	24.07	16.93			24.31
120	-	23.98	17.35			24.93
125	-	24.30	17.64			25.55
130	-	24.15	17.62			25.79
135	-	24.18		23.70		26.21
140	-	24.42		25.12		26.11
145	-	23.65		25.16		26.33
150	-	24.39		24.93		26.23
155	-	24.46		24.68		26.21
160	-	24.11		24.29		26.45
165	-	24.12		24.25		26.53
170	-	-	-	-		26.40
175	-	-	-	-		26.17
180	-	-	-	-		26.23
185	-	-	-	-		26.21
190	-	-	-	-		26.16
195	-	-	-	-		26.21
average	22.22	25.28	17.96	25.77	18.59	26.83

Table 8: Performance of Tag-less Back-translation compared to the baseline and standard back-translation models. BLEU Scores for each Checkpoint of the Models for German-English NMT (best single-checkpoint and average scores are shown in **bold**).

training step (thousands)	<i>baseline</i>	<i>standard_bt</i>	<i>tag-less_bt</i>		<i>tag-less_bt (joint BPE)</i>	
			pre-train	fine-tune	pre-train	fine-tune
5	2.82	2.13	0.43		0.37	
10	5.00	3.13	0.83		1.11	
15	6.30	4.45	1.09		2.02	
20	7.69	5.48	1.44		2.82	
25	9.13	6.30	1.66		3.35	
30	10.20	6.56	1.84		3.80	
35	11.23	7.76	2.12		4.37	
40	11.45	8.61	2.28		4.81	
45	11.84	9.76	2.38		5.05	
50	13.82	12.93	2.49		5.14	
55	15.73	16.24	2.77		5.41	
60	18.71	18.84	2.87		5.43	
65	19.74	20.21	2.85		5.38	
70	19.95	20.89	2.93			21.32
75	20.30	21.75	3.01			23.29
80	20.15	22.54	2.99			24.86
85	19.96	23.16		9.84		25.54
90	19.75	23.44		12.74		26.32
95	19.53	23.64		21.63		26.70
100	19.51	24.13		23.59		26.85
105	-	24.07		24.63		27.38
110	-	23.94		24.84		27.67
115	-	24.69		25.32		27.82
120	-	24.31		25.16		28.06
125	-	24.29		25.00		27.91
130	-	24.59		24.72		28.01
135	-	24.93		24.26		28.03
140	-	24.70	-	-		27.98
145	-	25.01	-	-		28.28
150	-	25.11	-	-		28.26
155	-	25.09	-	-		28.31
160	-	24.94	-	-		28.24
165	-	24.90	-	-		28.15
170	-	24.62	-	-		28.22
175	-	24.86	-	-	-	-
average	20.95	25.87	3.03	26.03	5.13	28.83

Table 9: Tagged Vs Tag-less Back-translation. We used the joint BPE for implementing the tag-less approach. BLEU Scores for each Checkpoint of the Models (best single-checkpoint and average scores are shown in **bold**).

training step (thousands)	Vi-En			De-En		
	<i>tagged_bt</i>	<i>tag-less_bt</i>		<i>tagged_bt</i>	<i>tag-less_bt</i>	
		pre-train	fine-tune		pre-train	fine-tune
5	1.78	2.87		2.86	0.37	
10	3.64	9.78		6.52	1.11	
15	9.88	13.06		10.55	2.02	
20	15.29	14.11		14.74	2.82	
25	16.99	15.03		16.48	3.35	
30	19.22	15.92		18.12	3.80	
35	19.62	15.95		19.82	4.37	
40	19.71	16.39		20.49	4.81	
45	20.99	17.08		21.47	5.05	
50	22.03	17.15		21.88	5.14	
55	21.85	17.29		22.29	5.41	
60	22.03	17.15		22.75	5.43	
65	22.33	17.37		23.19	5.38	
70	22.95	17.51		23.67		21.32
75	23.05	17.92		23.86		23.29
80	22.90	17.96		24.32		24.86
85	23.25	17.94		24.60		25.54
90	23.91	18.58		24.35		26.32
95	23.87	18.28		24.89		26.70
100	23.69	18.23		25.12		26.85
105	24.21	18.60		25.25		27.38
110	23.91	18.53		25.41		27.67
115	23.90		24.31	25.48		27.82
120	24.33		24.93	25.76		28.06
125	24.08		25.55	25.89		27.91
130	24.34		25.79	26.04		28.01
135	24.17		26.21	26.17		28.03
140	24.27		26.11	26.13		27.98
145	24.28		26.33	26.02		28.28
150	24.44		26.23	26.56		28.26
155	24.15		26.21	26.35		28.31
160	24.50		26.45	26.30		28.24
165	24.76		26.53	26.76		28.15
170	24.56		26.40	26.84		28.22
175	24.29		26.17	26.65	-	-
180	24.09		26.23	26.99	-	-
185	24.34		26.21	27.00	-	-
190	24.27		26.16	26.90	-	-
195	24.32		26.21	27.05	-	-
200	-	-	-	27.17	-	-
205	-	-	-	27.49	-	-
210	-	-	-	27.29	-	-
215	-	-	-	27.10	-	-
220	-	-	-	27.18	-	-
225	-	-	-	27.01	-	-
230	-	-	-	27.26	-	-
235	-	-	-	27.44	-	-
avarage	25.05	18.59	26.83	27.75	5.13	28.83

Table 10: Pre-training on the authentic data and fine-tuning on the synthetic data for Vietnamese-English NMT. BLEU Scores for each Checkpoint of the Models

training step (thousands)	<i>pre-train</i>	<i>fine-tune</i>
5	2.50	-
10	4.25	-
15	5.54	-
20	6.68	-
25	10.87	-
30	16.89	-
35	19.47	-
40	20.37	-
45	21.09	-
50	20.13	-
55	21.09	-
60	21.11	-
65	21.19	-
70	20.17	-
75	19.41	20.00
80	-	17.60
85	-	17.57
90	-	18.16
95	-	17.62
100	-	18.95
105	-	17.98
110	-	18.04
115	-	18.63
120	-	18.36
125	-	18.20
130	-	18.25
135	-	18.46
140	-	18.09
145	-	18.32
150	-	17.94
155	-	18.21
160	-	18.25
165	-	18.38
170	-	17.69
average	-	18.91

Table 11: Using different ratios of authentic to synthetic parallel data and its effect on the performance of Vietnamese-English NMT. Evaluation scores (BLEU) on the test set for each checkpoint (*tag-less_bt* colour code: BLACK – pre-train, RED – fine-tune)

training step (thousands)	<i>tagged_bt</i>			<i>tag-less_bt (joint BPE)</i>		
	1:1	1:3	1:5	1:1	1:3	1:5
5	2.31	1.78	1.87	1.65	2.87	1.82
10	3.31	3.64	3.23	8.68	9.78	6.16
15	7.35	9.88	4.05	11.93	13.06	12.81
20	14.45	15.29	5.08	13.52	14.11	14.20
25	17.12	16.99	5.71	14.36	15.03	13.63
30	18.54	19.22	13.36	15.24	15.92	16.08
35	20.45	19.62	16.41	15.49	15.95	15.24
40	21.76	19.71	17.28	15.90	16.39	16.34
45	21.86	20.99	19.20	16.18	17.08	17.32
50	21.96	22.03	19.93	16.05	17.15	16.36
55	22.87	21.85	20.33	16.52	17.29	17.45
60	22.71	22.03	20.85	16.58	17.15	16.99
65	23.05	22.33	21.62	16.30	17.37	17.83
70	23.33	22.95	21.83	16.82	17.51	16.94
75	22.84	23.05	21.84	16.80	17.92	18.36
80	23.25	22.90	23.05	16.76	17.96	17.86
85	23.51	23.25	22.56	16.54	17.94	18.14
90	22.99	23.91	23.15	16.47	18.58	18.44
95	23.06	23.87	23.62	16.69	18.28	18.52
100	22.89	23.69	23.37	23.04	18.23	18.31
105	22.96	24.21	23.58	24.05	18.60	18.80
110	22.80	23.91	23.49	24.81	18.53	18.42
115	22.92	23.90	23.90	24.87	24.31	19.22
120	22.83	24.33	23.69	25.50	24.93	19.07
125	22.29	24.08	23.79	24.86	25.55	18.47
130	22.53	24.34	24.10	26.29	25.79	19.44
135	22.43	24.17	24.01	25.77	26.21	19.48
140	22.36	24.27	24.52	25.95	26.11	19.16
145	22.34	24.28	24.57	25.94	26.33	19.14
150	22.08	24.44	23.88	25.96	26.23	19.01
155	22.08	24.15	24.62	25.76	26.21	24.82
160	21.74	24.50	24.49	-	26.45	25.79
165	22.14	24.76	24.12	-	26.53	26.21
170	22.73	24.56	25.47	-	26.40	26.27
175	-	24.29	-	-	26.17	26.51
180	-	24.09	-	-	26.23	26.78
185	-	24.34	-	-	26.21	26.91
190	-	24.27	-	-	26.16	26.60
195	-	24.32	-	-	26.21	26.57
200	-	-	-	-	-	26.58
205	-	-	-	-	-	26.64
210	-	-	-	-	-	25.70
215	-	-	-	-	-	26.29
average	22.73	25.05	25.47	26.15	26.83	27.14

Table 12: Fine-tuning the tagged and standard back-translations

training step (thousands)	<i>tagged.bt (1:5)</i>	<i>standard.bt (1:3)</i>
165	23.38	24.21
170	25.11	24.61
175	25.55	24.87
180	24.99	25.15
185	24.79	24.11
190	23.58	23.98
195	23.37	23.36
200	23.15	22.97
average	25.64	25.32

Table 13: This table shows how often a conclusion with 95% statistical significance is made for comparing the various approaches. We used different sample sizes of 100, 500 and 1000 sentences for each of the approach on English-Vietnamese and English-German low resource NMT.

System Comparison	Vi-En				De-En			
	BLEU difference	Sample size			BLEU difference	Sample size		
		100	500	1000		100	500	1000
Standard BT is better than baseline	3.06	100%	100%	100%	4.92	100%	100%	100%
Tagged BT is better than baseline	2.83	100%	100%	100%	6.81	100%	100%	100%
Tagged BT is better than Standard BT	-0.23	27%	28%	30.2%	1.89	100%	100%	100%
Tag-less BT is better than baseline	3.55	100%	100%	100%	5.09	100%	100%	100%
Tag-less BT better than Standard BT	0.49	92.8%	94.6%	94.8%	0.16	87.1%	87.4%	84%
Tag-less BT is better than Tagged BT	0.72	98.8%	98.2%	98%	-1.72	0%	0%	0%
Tag-less BT (joint BPE) is better than baseline	4.61	100%	100%	100%	7.89	100%	100%	100%
Tag-less BT (joint BPE) is better than Standard BT	1.55	100%	100%	100%	2.96	100%	100%	100%
Tag-less BT (joint BPE) is better than Tagged BT	1.78	100%	100%	100%	1.08	100%	100%	100%

[Click here to view linked References](#)

Noname manuscript No.
(will be inserted by the editor)

Tag-less Back-Translation

Idris Abdulmumin · Bashir Shehu
Galadanci · Aliyu Garba

This work is supported by the National Information Technology Development Agency under the National Information Technology Development Fund PhD Scholarship Scheme 2018

Received: date / Accepted: date

Abstract An effective method to generate a large number of parallel sentences for training improved neural machine translation (NMT) systems is the use of the back-translations of the target-side monolingual data. The standard back-translation method has been shown to be unable to efficiently utilize the available huge amount of existing monolingual data because of the inability of translation models to differentiate between the authentic and synthetic parallel data during training. Tagging, or using gates, has been used to enable translation models to distinguish between synthetic and authentic data, improving standard back-translation and also enabling the use of iterative back-translation on language pairs that underperformed using standard back-translation. In this work, we approach back-translation as a domain adaptation problem, eliminating the need for explicit tagging. In the approach – *tag-less back-translation* – the synthetic and authentic parallel data are treated as out-of-domain and in-domain data respectively and, through pre-training and fine-tuning, the translation model is shown to be able to learn more efficiently from them during training. Experimental results have shown that the approach outperforms the standard and tagged back-translation approaches on low resource English-Vietnamese and English-German neural machine translation.

Keywords tagged back-translation · tag-less back-translation · neural machine translation · machine translation · natural language processing

Idris Abdulmumin
Computer Science, Ahmadu Bello University, Zaria, Kaduna, Nigeria.
Tel.: +234-806-295-5509
E-mail: iabdulmumin@abu.edu.ng

Bashir Shehu Galadanci
Software Engineering, Bayero University, Kano, Kano, Nigeria.
E-mail: bsgaladanci.se@buk.edu.ng

Aliyu Garba
Computer Science, Ahmadu Bello University, Zaria, Kaduna, Nigeria.
E-mail: algarba@abu.edu.ng

1 Introduction

Neural Machine Translation (NMT) [?, ?, ?] has been the state-of-the-art approach for machine translation in recent years [?, ?], outperforming Phrase-Based Statistical Machine Translation [?] when qualitative parallel data between the languages is available in abundance [?]. This training dataset is usually scarce and expensive to compile for many language pairs. Recently, researchers have proposed methods to exploit the easier-to-get monolingual data of one or both of the languages to augment the available parallel data and improve the performance of the translation models. Such methods include integrating a language model [?], back-translation [?, ?, ?], forward translation [?] and dual learning [?]. The back-translation approach is simple and has been the most effective technique yet for NMT [?, ?]. The method involves training a target-to-source (backward) model on the available authentic bitext. The backward model is then used to translate a large amount of monolingual sentences in the target language into synthetic source sentences, generating the synthetic parallel data. The authentic and synthetic parallel data are then mixed to train a source-to-target (forward) model.

It has been shown that as the amount of monolingual data used in back-translation continues to increase, a point is reached when the model stops learning useful representation and, therefore, the performance of the model starts to drop. This is because the usually noise-infested synthetic data starts to overwhelm the authentic data and the model starts to completely unlearn the correct parameters it learns from the authentic training data [?]. Extensive studies by [?] have shown that in low resource NMT, noising beam search outputs improve the models more than other generation methods such as sampling. The authors claimed that the method enhances source-side diversity. But the works of [?, ?] found that the noising technique is only a form of tagging, indicating to the model that the noised data is back-translated, enabling it to treat the synthetic data as belonging to a different domain. The model then learns different representations, optimally, from the two data. They, instead, introduced the use of explicit tags (and gates) to indicate synthetic inputs. The tagging approach was shown to outperform the standard back-translation.

In this work, we approach back-translation as a domain adaptation problem, simplifying the works of [?, ?, ?] that explicitly differentiate between the two data using noise/tags/gates. Instead of tagging the synthetic data, our approach – the *tag-less back-translation* – aims to enable the model to learn efficiently from the two data through pre-training and fine-tuning. Instead of relying on the model to differentiate between the data, we used the synthetic data as generic domain (out-of-domain) and pre-train the model on this data. We then used the authentic data as in-domain to fine-tune the pre-trained translation model. We hypothesize that although the tagging and noising approaches improve the forward models, our domain-adaptation-tailored approach will provide a flexible method of maximizing the gains in the quantity of the synthetic data and efficiently utilizing the quality in the authentic parallel data. The approach will also enable the use of different training set-

tings on the different data, as obtainable in domain adaptation strategies. It also gets better as more research and more improved ways of domain adaption are proposed.

In domain adaptation, the generic model is not always expected to perform very well in the domain it is to be deployed, hence the model is fine-tuned with a usually smaller but in-domain data. In many languages, the in-domain data is usually low-resourced or non-existent: having the same issue as in low resource neural machine translation. The in-domain data in itself is not sufficient to create a good model while the more abundant out-of-domain data performs poorly when deployed in the target domain. Mixing the two data results in the in-domain data to be lost in the out-of-domain data and the resulting model is not able to also perform well in the target domain. The larger out-of-domain data is, therefore, used to pre-train a model and the weights of this model are used to initialize the training of the in-domain translation model – a technique referred to as *fine-tuning* [?]. When a different language pair is used for pre-training than that used during fine-tuning, the approach is regarded to as transfer learning [?,?].

We make the following contributions in this paper:

- we proposed a novel approach that enables a translation model that is trained on synthetic and authentic parallel data to be able to efficiently learn from the the two data, utilizing the different advantages presented by each.
- we successfully applied pre-training and fine-tuning to enable the forward model in back-translation to differentiate between synthetic and authentic data during training, achieving a superior performance to standard and the successful tagged back-translation approaches,
- experimental results have shown that the approach is superior to the standard and tagging back-translation approaches in low resource English-Vietnamese and English-German neural machine translation systems.

The remaining sections are as follows: Section ?? reviews relevant literature on NMT, leveraging monolingual data in NMT and pre-training and fine-tuning. Section ?? explains the tag-less back-translation approach, Section ?? describes the data and experimental set-up used in training the models, Section ?? discusses the results obtained after the experiment. We discuss further the findings in Section ?. Finally, in Section ?, we concluded the work and suggest future directions.

2 Related Works

This section presents prior work on NMT, back-translation and pre-training in NMT.

2.1 Neural Machine Translation (NMT)

The NMT is based on a sequence-to-sequence encoder-decoder system with attention mechanism [?, ?, ?]. The encoders and decoders are made of neural networks that model the conditional probability of a target sentence y given the source sentence x : $p(y|x)$. The encoder converts the input in the source language into a set of vectors while the decoder converts the set of vectors into the target language through an attention mechanism, one word at a time. The attention mechanism was introduced to keep track of context in longer sentences [?].

The NMT model produces the translation sentence by generating one target word at every time step. Given an input sequence $X = (x_1, \dots, x_{T_x})$ and previously translated words $(y_1, \dots, y_{i-1})$, the probability of the next word y_i is

$$p(y_i|y_1, \dots, y_{i-1}, X) = g(y_{i-1}, s_i, c_i) \quad (1)$$

where s_i is the decoder hidden state for time step i and is computed as

$$s_i = f(s_{i-1}, y_{i-1}, c_i). \quad (2)$$

Here, f and g are nonlinear transform functions, which can be implemented as long short-term memory (LSTM) network [?] or gated recurrent units (GRU) [?] in recurrent neural machine translation (RNMT), and c_i is a distinct context vector at time step i , which is calculated as a weighted sum of the input annotations h_j

$$\sum_{j=1}^{T_x} a_{i,j} h_j \quad (3)$$

where h_j is the annotation of x_j calculated by a bidirectional Recurrent Neural Network. The weight $a_{i,j}$ for h_j is calculated as

$$a_{i,j} = \frac{\exp e_{i,j}}{\sum_{t=1}^{T_x} \exp e_{i,t}} \quad (4)$$

and

$$e_{i,j} = v_a \tanh(W s_{i-1} + U h_j) \quad (5)$$

where v_a is the weight vector, W and U are the weight matrices.

All of the parameters in the NMT model, represented as θ , are optimized to maximize the following conditional log-likelihood of the M sentence aligned bilingual samples

$$L(\theta) = \frac{1}{M} \sum_{m=1}^M \sum_{i=1}^{T_y} \log(p(y_i^m | y_{<i}^m, X^m, \theta)) \quad (6)$$

To remove the recurrence and enable parallelization across multiple GPUs during training, the convolutional neural networks were used to create the convolutional NMT (CNMT) encoder-decoder architecture [?, ?]. The CNMT

utilizes 1-dimensional convolutional layers followed by gated linear units, GLU [?]. The decoders compute and apply attention to each of the layers. The model uses positional embeddings along with residual connections [?].

The transformer [?,?] architecture was introduced to remove the recurrence and convolutions of previous architectures. The transformer is based solely on multi-headed self-attention layers. It enables parallelization across multiple GPUs, thereby, reducing training time. The architecture is used in current state-of-the-art translation systems [?,?].

In this work, we used a unidirectional LSTM encoder-decoder architecture with Luong attention [?]. This is a simple recurrent neural network RNMT architecture. Our approach is not architecture-dependent and can be applied to the other architectures or other more enhanced implementations of the RNMT.

2.2 Leveraging Monolingual Data for NMT

The use of monolingual data of the target and/or source language has been studied extensively to improve the performance of neural translation models, especially in low resource settings. [?] explored integrating language models trained on monolingual data into NMT systems, [?,?] proposed augmenting a copy or slightly modified copy respectively of the target data as source, [?] proposed the back-translation approach, [?] proposed the forward translation and [?] used both forward and back-translations to improve the translation models. The back-translation approach has been shown to outperform other approaches in low and high resource languages [?,?].

Various studies have investigated back-translation to improve the backward model, to select the most suitable generation/decoding methods for generating the synthetic data and to reduce the impact of higher ratio of the synthetic to the authentic bitext. The quality of the models trained using back-translation depends on the quality of the backward model [?,?,?,?,?]. To improve the quality of the synthetic parallel data, [?] used iterative back-translation – iteratively using the back-translated data to improve both the backward and forward models. [?] and [?] used high resource languages through transfer learning and [?] explored the use of both target and source monolingual data to improve both the backward and forward models. [?] trained a bilingual system based on [?] to do both forward and backward translations, eliminating the need for two separate models. [?] used Transductive data selection methods to select monolingual data that are in the same domain as the test set for back-translation, improving performance.

The works of [?,?] have found that the ratio of synthetic to authentic data affects the performance of the models most. When the ratio is high, the model tends to learn more from the synthetic data, which contains more mistakes than the authentic data. Investigations have found that the sampling approach of synthetic data generation and adding noise to beam search output outperforms the regular beam decoding technique [?,?]. These approaches were

said to improve the models by enhancing source-side diversity. [?] claimed, instead, that the noise only indicates to the model that the input is either synthetic or authentic, enabling the model to better utilize the two data. [?] and [?] used tags (and gates) to enable the model to distinguish between the data and the approach has been shown to efficiently utilize more synthetic data, outperforming standard back-translation and enhancing the efficiency of iterative back-translation.

2.3 Domain Adaptation

Domain adaptation is the use of a usually few amount of in-domain data to improve the performance of an out-of-domain (general purpose) model before deployment. The amount of the in-domain parallel data is usually not sufficient to train a very good model and the general purpose models usually performs poorly [?]. There are two categories of domain adaptation – *data centric* and *model centric* [?] with each having several techniques. The techniques in these classifications include using monolingual data [?], synthetic data generation [?], using data selection [?] and using tagged out-of-domain parallel data [?] and fine-tuning [?]

Pre-training has been used successfully in various machine learning tasks to improve performance when the data is not enough to train a good enough model. It was used for training word embeddings [?], in computer vision [?], fine-tuning NMT models [?] and as transfer learning in low resource NMT [?,?]. The transfer learning for machine translation approach involves training a model on a high resource language pair and transferring the training on a low resource pair. The works of [?,?,?] have shown tremendous improvements over models that are trained with the low resource data from scratch.

In back-translation, [?] showed that fine-tuning a pre-trained model on in-domain data improves the quality of back-translated model. [?] pre-trained the model on the authentic data and fine-tunes it on the mixed synthetic and authentic data. [?] and [?] pre-trained a model on a high resource language and fine-tunes it on a low resource language pair.

3 The Proposed Method

The approach is shown in Fig. 1. As illustrated in Algorithm 1, the authentic parallel data: $D^P = \{(x^{(u)}, y^{(u)})\}_{u=1}^U$ is used to train a target-to-source model, $M_{x \leftarrow y}$. This model – the *backward* model – is used to translate the monolingual target data, $Y = \{(y^{(v)})\}_{v=1}^V$, to generate the synthetic parallel data: $D' = \{(x^{(v)}, y^{(v)})\}_{v=1}^V$. Instead of mixing the two data to train a forward (target) model, we used only the synthetic data to pre-train the forward model, $M_{x \rightarrow y}$, until no improvement is observed on the development set. Finally, the forward model is fine-tuned on authentic data.

It was shown in [?] that using different vocabulary each during pre-training and fine-tuning leads to drop in performance because, it was said, independent

Fig. 1: Tag-less Back-Translation: Training the forward model on the synthetic parallel data generated using the backward model. The forward model is then fine-tuned on the authentic data.

vocabulary use different identifiers even for the same subwords and the network loses benefits of the weights learned during pre-training. The authors proposed learning a joint BPE on a mixture of both the pre-training and fine-tuning data and this has been shown to achieve better results in domain adaptation. In this approach, we have access to both the out-of-domain (synthetic) and the in-domain (authentic) parallel data. This, therefore, enables us to learn a joint BPE and build the training vocabulary for both pre-training and fine-tuning.

Algorithm 1: Tag-less Back-Translation

Input: Parallel data $D^P = \{(x^{(u)}, y^{(u)})\}_{u=1}^U$ and
Monolingual target data $Y = \{(y^{(v)})\}_{v=1}^V$

- 1: **procedure** BACK-TRANSLATION
- 2: Train backward model $M_{x \leftarrow y}$ on bilingual data D^P
- 3: Use $M_{x \leftarrow y}$ to create $D' = \{(x^{(v)}, y^{(v)})\}_{v=1}^V$, for $y \in Y$;
- 4: Pre-train forward model $M_{x \rightarrow y}$ on parallel data D' ;
- 5: Fine-tune the forward model $M_{x \rightarrow y}$ on parallel data D^P ;
- 6: **end procedure**

Output: forward model $M_{x \rightarrow y}$

Table 1: Data Used

data	train		dev	test
	sentences	words (vocab)		
		EnVi		
IWSLT'15 En-Vi	133, 317	2,837,240 (50,045)2,688,387 (103,796)	1, 553	1, 268
IWSLT'14 En-De	153, 348	2,706,255 (54,169)3,311,508 (25,615)	6, 970	6, 750
WMT'14 En-De – Monolingual English	666,585	16,700,106 (344,138)	-	-

4 Experiments

4.1 Set-up

We used the TensorFlow [?] implementation of the OpenNMT [?] framework to train the models – the NMTSmallV1 configuration. The set-up is based on the NMTSmallV1 configuration. Specifically, the configuration is a 2-layer unidirectional LSTM encoder-decoder model with Luong attention [?]. It has 512 hidden units and a vocabulary size of 50,000 for both source and target languages. We used Adam [?] optimizer and a batch size of 64 with a dropout probability of 0.3, a static learning rate of 0.0002 and the models are evaluated on the development set after every 5,000 training steps. The models were evaluated using the bi-lingual evaluation understudy metric, BLEU [?], specifically the multi-bleu [?] implementation. The models are trained until there is no improvement of over 0.2 BLEU after four training steps. As stated in Section ??, the learning of BPE on the training data and the building of training vocabulary for both pre-training and fine-tuning was done on the mixture of the synthetic and authentic parallel data. During fine-tuning, we, therefore, only change the training data.

4.2 Data

For this work, we use the preprocessed low resource English-Vietnamese parallel data [?] of the IWSLT 2015 Translation task [?]. We used the 2012 and 2013 test sets for development and testing respectively. We also used the data from the IWSLT 2014 German-English shared translation task [?] as the second language pair, pre-processed using the data clean-up as well as the train, development and test split in [?]. For the monolingual data, we used the pre-processed English monolingual data of WMT 2014 English-German translation task [?]. We shuffled the monolingual data and selected 666,585 monolingual sentences which is five times as much as the En-Vi parallel data. The statistics

of the datasets are shown in Table ?? . We learned byte pair encoding (BPE) [?] with 10,000 merge operations on the training dataset and applied it on the train, development and test datasets. Afterwards, we build the vocabulary on the training dataset. For all the experiments, we used thrice as much of the monolingual as the available parallel data in both of the languages except when we experimented with the ratio of 1:5 (parallel to monolingual data) for the English-Vietnamese NMT.

4.3 Models

To compare the performance of our approach with that of the previous works, we implemented the following methods to train translation models on the English-Vietnamese and English-German NMT using the data presented in Section ?? above. All models were trained using the same settings stated in Section ??

- We first train *baseline* models using the available authentic parallel data only. In the models, the baselines have the English language as the target language – Vi-En and De-En.
- We then train the backward models also on the authentic parallel data using English language as the source language – En-Vi and En-De. The models are used for the generation of the additional synthetic parallel data for the back-translation approach.
- We implemented the various back-translation strategies namely standard back-translation – *standard_bt*, the tagged back-translation – *tagged_bt* and the tag-less back-translation – *tag-less_bt* (*joint BPE*) using the authentic and synthetic parallel data.

5 Results

All scores reported are statistically significant with $p < 0.05$. We used the paired bootstrap resampling of [?] as implemented in [?] to estimate the statistical significance confidence scores. See Table ?? in Appendix 1 for confidence scores.

5.1 English-Vietnamese Low Resource NMT

The evaluation scores of the best models and the improved models obtained after taking the checkpoint averaging of the last 8 checkpoints are shown in Table ?? . We first created a forward Vietnamese-English (Vi-En) model, *baseline*, on the available authentic parallel data. The model trained for 75,000 steps before the stopping condition was met. The baseline model was trained further to 110,000 training steps but the performance continued to flatten without observing any improvement. The model achieved the best single-checkpoint

Table 2: Performance of the Tag-less Back-translated model compared to the baseline and standard back-translation models for Vietnamese-English and German-English translations. Evaluation scores on the test set. The tag-less approaches show results of pre-training and fine-tuning.

data		<i>baseline</i>	<i>standard_bt</i>	<i>tag-less_bt</i>	<i>tag-less_bt</i> (joint BPE)
Vi-En	best checkpoint				
	BLEU (training step)	21.19 (65k)	24.46 (155k)	17.85 (90k) 25.16 (145k)	18.60 (105k) 26.53 (165k)
	average	22.22	25.28	17.96 25.77	18.59 26.83
De-En	best checkpoint				
	BLEU (training step)	20.30 (75k)	25.11 (150k)	3.01 (75k) 25.16 (120k)	5.43 (60k) 28.31 (155k)
	average	20.95	25.87	3.03 26.03	5.13 28.83

score of 21.19 BLEU at the 65,000th. We then trained a backward (En-Vi) model, *backward*. After the stoppage condition were met, after 55,000 training steps, the best performing single-checkpoint for the backward model achieved a BLEU score of 24.78 at the 50,000th training step. Averaging the last 8 checkpoints gave the best performance – 25.79 BLEU. This average model was used to back-translate the monolingual English data to generate synthetic parallel data. We mixed the two data – synthetic and authentic – without differentiating between the two and used the resulting large dataset to train a forward model. We labelled this model as *standard_bt* – for standard back-translation [?]. This model was trained for 165,000 before the stopping condition were met. It achieved a single-checkpoint best BLEU score of 24.46 at the 155,000th training step.

We mixed the synthetic and authentic parallel data and learned a joint BPE on the resulting training dataset and build the training vocabulary. We applied the BPE on the synthetic data for pre-training and on the authentic data for fine-tuning. We trained a model, labelled *tag-less_bt (joint BPE)*, using this approach. The model achieved a single-checkpoint score of 18.60 BLEU during pre-training and improved to 26.53 BLEU after fine-tuning. The average fine-tuned model was better by about 0.30 BLEU. The average pre-trained model performed very low compared to the baseline and the standard back-translation models – 18.59 BLEU vs 22.22 and 25.28 BLEUs respectively. This is obviously because the quality of the data used in the training the model – the synthetic data – is lower than that of the other two. The quality of the synthetic data, although generated from a reasonably good backward model, is still not sufficient to train a model whose quality can compare to the other models that are trained in whole or in part with the authentic data. Fine-tuning the model on the authentic data results in a sharp rise in performance.

Fig. 2: Tag-less back-translation: pre-training on synthetic data and fine-tuning on authentic data. Showing how this technique compares to the baseline and the standard back-translation approaches on the test set.

The model was fine-tuned until the stopping condition was met. The approach outperformed the baseline and standard back translation models by 5.34 and 2.07 BLEUs respectively. The gap in performance was, however, reduced to 4.61 and 1.55 BLEUs after checkpoint averaging.

We experimented the other pre-train and fine-tune approach, learning the BPE only on the synthetic data. We build the vocabulary on the synthetic training data after applying the BPE. The synthetic corpus was used to pre-train a forward model for 130,000 steps, achieving a single-checkpoint best score of 17.85 BLEU. The authentic parallel data was then used to fine-tune the model for a further 35,000 training steps. Stopping at each of these steps were based on the stopping condition. The performance of the *tag-less_bt* model improved to 25.16 BLEU after fine-tuning. Although this approach was shown to outperform the baseline and standard back-translation, it underperformed the joint BPE implementation of the tag-less approach by 1.06 BLEU. In Figures ??a and ??b, we show how the BLEU scores continue to improve with increase in the training steps. The model trained using the *tag-less_bt (joint BPE)* approach continued to outperform the three others after fine-tuning.

5.2 English-German Low Resource NMT

We conducted the same set of experiments presented in section ?? on the second low resource dataset, the English-German IWSLT’14 parallel dataset. This data, as presented in Table ??, is made up of a little bit more than 150,000 parallel sentences. We first trained a backward (En-De) model on the available parallel data. This model maxed-out performance on the test set, based on the set-up, at 10.25 BLEU after averaging the last 8 checkpoints. It stopped training at the 80,000th training steps and achieving the best single model performance at the 65,000th – 10.03 BLEU. We used the average model to generate the synthetic data, translating the available English monolingual

Fig. 3: Tagged Vs Tag-less back-translation.

data. We trained four separate forward (De-En) models based on the approaches we explained earlier. The first is the baseline trained on the available authentic data, the *standard.bt* on the mixture of the authentic and synthetic data without differentiating, the *tag-less.bt* pre-trained on the synthetic data and fine-tuned on the authentic data having learned the BPE on the synthetic data and updating the vocabulary before fine-tuning and, finally, the *tag-less.bt (joint BPE)* trained also using the tag-less approach but having learned the BPE and built the vocabulary on the mixture of the synthetic and authentic data.

The results of evaluating the models after training using the various approaches are presented in Table ?? The baseline achieved a modest average performance of 20.95 BLEU after training for 100,000 training steps. The performance on the dataset was improved after applying standard back-translation, achieving a huge +4.92 BLEU improvement over the baseline. The tag-less approach, though better, did not achieve a huge improvement over back-translation (only +0.16 BLEU) but after applying the improved tag-less (joint BPE), as shown in the previous section, we achieved huge +2.96 BLEU increase in performance. This +2.8 and +7.88 BLEUs over the previous tag-less approach and the baseline respectively.

For all the subsequent experiments, unless stated otherwise, we used the joint BPE technique to implement the tag-less back-translation approach as it is shown to be the most successful variant.

5.3 Tagged Vs Tag-Less Back-translation

We compared the performance of the *tag-less.bt* model – our technique – with that of the successful tagged back-translation of [?] on the English-Vietnamese data. The synthetic sources were labelled with the $\langle BT \rangle$ token at the beginning of each sentence and mixed with the authentic sources to generate the mixed tagged parallel corpus. This mixed data is used to train the forward

Table 3: Performance of Tag-less Back-translation compared to the Tagged back-translation model for Vietnamese-English translation. Evaluation scores on the test set.

		<i>tagged_bt</i>	<i>tag-less_bt (joint BPE)</i>
En-Vi	best checkpoint BLEU (training step)	24.76 (165k)	26.53 (165k)
	average	25.05	26.83
En-De	best checkpoint BLEU (training step)	27.49 (205k)	28.31 (155k)
	average	27.75	28.83

Table 4: Performance of Tag-less Back-translation on the test set: pre-training on synthetic data and fine-tuning on authentic data Vs pre-training on authentic data and fine-tuning on synthetic data for Vietnamese-English.

		<i>tag-less_bt</i>	<i>reverse_tag-less_bt</i>
best checkpoint BLEU (training step)		25.16 (145k)	21.19 (65k) - pre-train 18.95 (100k) - fine-tune
	average	25.77	18.91

tagged back-translation model – *tagged_bt*. The *tagged_bt* model stopped at 125,000 steps and the training was continued up to 195,000 steps to equal the number of training steps reached by the *tag-less_bt* model. While the tagged approach underperformed the best score of our technique by 1.78 BLEU, it was able to outperform the single-checkpoint standard back-translation by 24.76 to 24.46 BLEUs respectively (+0.3 BLEU) but underperformed the average standard back-translation model by 0.23 BLEU.

Finally, we trained a forward model using the tagged back-translation for English-German NMT to compare the performance with our approach on this data. The tagged approach took a further 50,000 training steps to reach a single model best of 27.49 BLEU, but still underperforming the tag-less approach by 0.82 BLEU. The best model obtained after averaging checkpoints was also achieved using our approach, a performance of 28.83 BLEU compared to the tagged 27.75 BLEU, an improvement of 1.08 BLEU. The performances of these models, evaluated on the test set is shown in Table ???. On this data, the tagged approach performed better than the standard back-translation by +1.88 BLEU on the average models. It can be seen in Figures ?? and ?? that in both of the experiments conducted on the two data, our tag-less approach out-performed the rest of the back-translation approaches.

This supports the hypothesis that although the tagged back-translation involves explicit differentiating between the two data using tags, the model trained on the approach may not be able to differentiate between them com-

Fig. 4: Fine-tuning the baseline model on the synthetic data. Evaluation scores on the test set.

pletely during training as observed in the mixed performance of the models trained on the two different data.

5.4 Fine-tuning: Synthetic Vs Authentic Data

Our technique proposed pre-training the forward model on the synthetic parallel data and fine-tuning the model afterwards on the authentic data. This was proposed to enable the model to unlearn the mistakes it learned from the synthetic data using correct sentences in the authentic parallel data. We experiment the other way round to investigate the effects of pre-training on the authentic data and fine-tuning on the synthetic data. We used the baseline as the pre-trained model and fine-tune it on the synthetic data. This approach was labelled as *reverse_tag-less_bt*. This approach did not show any benefit to the final forward model, see Fig. 3. As expected, the performance of the model decreased and the curve flattens as the number of training steps increases. The best and average scores are shown in Table 3.

5.5 Quantity of Monolingual Data

As stated earlier, it was found that as the more synthetic data increases, a point is reached where the performance starts to deteriorate [?]. Instead, our work hypothesizes that the performance of the model will start to decrease only if it is not able to differentiate between the synthetic and authentic training data and, therefore, efficiently learning from the two. We also pointed out

Table 5: Using different ratios of the authentic to synthetic data for Vietnamese-English translation. Evaluation scores of the models on the test set.

	<i>tagged.bt</i>			<i>tag-less.bt</i> (joint BPE)		
	1:1	1:3	1:5	1:1	1:3	1:5
best checkpoint BLEU (training step)	22.73 (85k)	24.76 (165k)	24.62 (155k)	26.29 (130k)	26.53 (165k)	26.91 (185k)
average	22.73	25.05	25.47	26.15	26.83	27.14

that since the data is mixed in both the standard and tagged back-translation approaches, the model may not be able to completely differentiate between the data, although in the latter approach, the model is expected to treat the tagged synthetic sources as a different domain. We, therefore, experiment with different ratios of the authentic to synthetic data to verify this claim. We sample the authentic to synthetic data in the ratios, 1:1, 1:3 and 1:5. The results are shown in Table ??.

In the tagged approach, the single-checkpoint best scores continue to rise from using the same amount of monolingual data for back-translation to using three times the authentic data of the monolingual data for back-translation. But, as observed, the performance dropped slightly when we used five times the amount of available parallel data. However, the performance of the tag-less back-translation models continues to increase steadily when the ratio of authentic to monolingual data is increased. We observed the performance to improve by about 0.25 BLEU when the amount of synthetic data is tripled and doubled to about 0.5 BLEU after adding another double amount of the synthetic data to the training data. It can also be observed that there was a very low improvement over the performance of *baseline* and serious underperformance compared to the tag-less approach when we used the same amount of synthetic data with the authentic data to train the models – 22.73 BLEU vs 22.22 and 26.15 BLEUs respectively. Overall, we obtained a 3.42, 1.78 and 1.67 BLEU improvements on the average models using the tag-less approach over the tagged approach on the ratios experimented respectively.

It can be observed also that the performance of the model trained on the 1:1 ratio of monolingual to synthetic data using our approach is very good compared to the model trained using the same amount of data in the tagged approach and subsequent increase in training data leads to steady improvements that at 1:5 ratio, the performance was improved by about 1 BLEU. This steady improvements can show that the model learned useful knowledge on the authentic data but only used the synthetic data for further improvements. Following this trend, we can, therefore, conclude that with more synthetic data compared to the authentic data, the model will only continue

Table 6: Before and after fine-tuning the English-Vietnamese standard and tagged back-translation NMT models on the authentic data. Evaluation scores on the test set.

	<i>tagged_bt (1:5)</i>		<i>standard_bt (1:3)</i>	
	before	after	before	after
best checkpoint BLEU (training step)	24.62 (155k)	25.55 (175k)	24.46 (155k)	25.15 (180k)
average	25.47	25.64	25.28	25.32

to learn and increase its performance if useful representations are learnable on the synthetic data.

5.6 Fine-tuning Standard And Tagged Back-Translations

The work of [?] reported no observable advantage of using the authentic data to train the forward model and fine-tuning it henceforth on the mixed data. Instead, we experimented training the forward model on the mixed data first and then fine-tune it on the authentic data. As shown in Table 6, this approach reaches the same performance as the old tag-less approach – 25.15 BLEU – using the same amount of synthetic sentences albeit after 30,000 more training steps but still underperforming the joint BPE tag-less back-translation’s 26.53 BLEU although training for additional 15,000 training steps. The better joint BPE tag-less approach converges earlier than fine-tuning the standard back-translation model, at 165,000 training steps.

We also explored the use of fine-tuning to determine whether or not the tagged approach will be able to cover the difference in performance with the tag-less approach. After fine-tuning the *tagged_bt (1:5)* model for a further 35,000 training steps, the performance gained was a significant +0.93 BLEU and only 0.17 BLEU over the average after just 20,000 steps of fine-tuning. The performance was still short of the *tag-less_bt (joint BPE) (1:5)* by a significant 1.36 BLEU.

6 Discussion

In this work, we proposed an approach for training the forward model in back-translation without using tags or noising the synthetic data. Translation models that are trained on the synthetic and authentic data have been shown to perform better when they are able to differentiate between the two data. Previous approaches have relied on the use of noise in back-translation [?] especially on low resource languages to improve the performance of models. The authors thought that the approach ensures source-side diversity which has been shown to benefit the models [?]. The approach was found out to

only indicate to the forward model that the noised data is synthetic, enabling it to treat the data differently from the authentic data [?]. The use of tags has been shown to improve the performance of such models. In this work, we eliminated the need of using of the tags and showed that although it was successful at improving the performance – proving it successful at indicating to the model that a data is synthetic and not authentic – domain adaptation methods are more capable of ensuring the model differentiate between the data. The ability for the model to separate between the data is even more important in low resource languages where the available data is not enough to train standard backward model, thus generating synthetic data with a lot of noises.

Domain adaptation techniques techniques in machine translation ensures that a better model is trained, leveraging on a larger parallel data of either the same language pair but in a different domain – fine-tuning – or a different language pair – transfer learning. In this technique, the two data are not tagged, mixed and left to the model to differentiate between them. They, rather, are used at the different stages of the training and this ensures the model performs in the target domain as expected. We utilized the synthetic data – which is bigger but more prone to translation noises – as the generic domain and the authentic data – smaller but having more quality – as the in-domain. This selection was not done until the reversed approach was shown not give the desired performance. The superiority of the approach over the successful tagging was shown through experimental results conducted on two low resource language pair: English-Vietnamese and English-German. In each of the languages considered, we obtained an improvement of more than 1 BLEU points over the tagged approach that outperformed the baseline and standard back-translation models.

We also test the performance of our technique when the amount of monolingual data is increased. We used different ratios of the authentic parallel to monolingual data used. We found that our technique was not only able to handle the increase in the synthetic data, but was able to attain rapid improvement given the smallest amount of synthetic data. We obtained a superior performance by a whopping 3.56 BLEU using the tag-less approach over the tagged approach when the amount of monolingual data is the same as the authentic data. The performance continued to steady increase as the amount of monolingual data is increased. The tagged approach could only handle tripling the amount of synthetic data but the performance started to decrease when the synthetic data was increased further. Using the same amount of synthetic data in ratios 1:1, 1:3 and 1:5, our technique outperformed the tagging technique by 3.42, 1.78 and 1.67 respectively (see average scores in Table ??).

Our approach also provides one with the flexibility of using state-of-the-art domain adaptation methods to improve the performance of the already successful back-translation approach. Techniques such as using different dropout and/or learning rate during pre-training and fine-tuning may improve the performance of the forward model. The method may also be applied in high-resource languages since both of these settings – low and high resource – can

benefit from the ability of the forward model to differentiate between synthetic and authentic data.

7 Conclusions and Future Work

This work has shown that an NMT model pre-trained on synthetic data and fine-tuned on the authentic data outperforms the rather successful method of tagging the synthetic data in low resource NMT by enabling the forward model to differentiate between the authentic and synthetic training data. The approach, however, does not improve the performance when it is reversed and the forward model is pre-trained on the authentic data and then fine-tuned on the synthetic data. As expected, the reverse approach makes the model to unlearn the useful representations learned in favour of the noise in the synthetic data. This justifies our hypothesis that without differentiating between the two data, the synthetic data is most likely to hurt the performance of the forward model.

It was shown also, in this work, that the more synthetic data used, the better the performance of the forward model, though the most effective ratio was not yet determined through thorough experimentation. This will inform the basis of future works. We experiment fine-tuning the models trained using the standard and tagged back-translation approaches. Experimental results showed the standard back-translation equalling the performance of a variant of the tag-less approach after many more rounds of training. The performance of the tagged approach improved considerably but still trailed the tag-less approach. The most successful of the tag-less approach has been the one that involves learning a joint BPE and building the training vocabulary on the mixture of the synthetic and authentic parallel data. This approach is made possible, unlike in other fine-tuning conditions, because both the generic (synthetic) data and the in-domain (authentic) data are available during the process.

For future work, the use of different settings – such as increasing or decreasing the learning rate, using dropout and L2 regularization, which may reduce overfitting on the in-domain (authentic) data as shown to be a likely problem in domain adaptation by [?] – for the pre-training and fine-tuning approaches can be explored to maximize the benefits of the domain adaptation approach in back-translation. The approach can also be investigated to improve the forward translation approach – which also leverages on the synthetic data for additional training data. Finally, we intend to investigate the technique in high resource languages in the future.

Conflict of interest

The authors declare that they have no conflict of interest.

Appendix

In this section, we provided the complete evaluation scores on the test set for all the models trained.

Table 7: Performance of Tag-less Back-translation compared to the baseline and standard back-translation models. BLEU Scores for each Checkpoint of the Models for Vietnamese-English NMT (best single-checkpoint and average scores are shown in **bold**).

training step (thousands)	<i>baseline</i>	<i>standard_bt</i>	<i>tag-less_bt</i>		<i>tag-less_bt (joint BPE)</i>	
			pre-train	fine-tune	pre-train	fine-tune
5	2.50	1.92	2.88		2.87	
10	4.25	3.50	7.54		9.78	
15	5.54	11.34	10.70		13.06	
20	6.68	15.83	12.86		14.11	
25	10.87	17.20	12.96		15.03	
30	16.89	19.47	14.07		15.92	
35	19.47	18.95	15.10		15.95	
40	20.37	20.81	15.59		16.39	
45	21.09	21.61	15.61		17.08	
50	20.13	22.18	15.99		17.15	
55	21.09	21.98	16.69		17.29	
60	21.11	22.35	16.66		17.15	
65	21.19	22.37	16.73		17.37	
70	20.17	23.43	17.63		17.51	
75	19.41	22.92	17.13		17.92	
80	20.43	23.41	17.33		17.96	
85	20.18	23.62	16.97		17.94	
90	19.80	23.71	17.85		18.58	
95	20.36	23.77	17.30		18.28	
100	19.48	23.73	17.31		18.23	
105	19.63	24.07	17.36		18.60	
110	19.77	23.99	17.34		18.53	
115	-	24.07	16.93			24.31
120	-	23.98	17.35			24.93
125	-	24.30	17.64			25.55
130	-	24.15	17.62			25.79
135	-	24.18		23.70		26.21
140	-	24.42		25.12		26.11
145	-	23.65		25.16		26.33
150	-	24.39		24.93		26.23
155	-	24.46		24.68		26.21
160	-	24.11		24.29		26.45
165	-	24.12		24.25		26.53
170	-	-	-	-		26.40
175	-	-	-	-		26.17
180	-	-	-	-		26.23
185	-	-	-	-		26.21
190	-	-	-	-		26.16
195	-	-	-	-		26.21
average	22.22	25.28	17.96	25.77	18.59	26.83

Table 8: Performance of Tag-less Back-translation compared to the baseline and standard back-translation models. BLEU Scores for each Checkpoint of the Models for German-English NMT (best single-checkpoint and average scores are shown in **bold**).

training step (thousands)	<i>baseline</i>	<i>standard_bt</i>	<i>tag-less_bt</i>		<i>tag-less_bt (joint BPE)</i>	
			pre-train	fine-tune	pre-train	fine-tune
5	2.82	2.13	0.43		0.37	
10	5.00	3.13	0.83		1.11	
15	6.30	4.45	1.09		2.02	
20	7.69	5.48	1.44		2.82	
25	9.13	6.30	1.66		3.35	
30	10.20	6.56	1.84		3.80	
35	11.23	7.76	2.12		4.37	
40	11.45	8.61	2.28		4.81	
45	11.84	9.76	2.38		5.05	
50	13.82	12.93	2.49		5.14	
55	15.73	16.24	2.77		5.41	
60	18.71	18.84	2.87		5.43	
65	19.74	20.21	2.85		5.38	
70	19.95	20.89	2.93			21.32
75	20.30	21.75	3.01			23.29
80	20.15	22.54	2.99			24.86
85	19.96	23.16		9.84		25.54
90	19.75	23.44		12.74		26.32
95	19.53	23.64		21.63		26.70
100	19.51	24.13		23.59		26.85
105	-	24.07		24.63		27.38
110	-	23.94		24.84		27.67
115	-	24.69		25.32		27.82
120	-	24.31		25.16		28.06
125	-	24.29		25.00		27.91
130	-	24.59		24.72		28.01
135	-	24.93		24.26		28.03
140	-	24.70	-	-		27.98
145	-	25.01	-	-		28.28
150	-	25.11	-	-		28.26
155	-	25.09	-	-		28.31
160	-	24.94	-	-		28.24
165	-	24.90	-	-		28.15
170	-	24.62	-	-		28.22
175	-	24.86	-	-	-	-
average	20.95	25.87	3.03	26.03	5.13	28.83

Table 9: Tagged Vs Tag-less Back-translation. We used the joint BPE for implementing the tag-less approach. BLEU Scores for each Checkpoint of the Models (best single-checkpoint and average scores are shown in **bold**).

training step (thousands)	Vi-En			De-En		
	<i>tagged_bt</i>	<i>tag-less_bt</i>		<i>tagged_bt</i>	<i>tag-less_bt</i>	
		pre-train	fine-tune		pre-train	fine-tune
5	1.78	2.87		2.86	0.37	
10	3.64	9.78		6.52	1.11	
15	9.88	13.06		10.55	2.02	
20	15.29	14.11		14.74	2.82	
25	16.99	15.03		16.48	3.35	
30	19.22	15.92		18.12	3.80	
35	19.62	15.95		19.82	4.37	
40	19.71	16.39		20.49	4.81	
45	20.99	17.08		21.47	5.05	
50	22.03	17.15		21.88	5.14	
55	21.85	17.29		22.29	5.41	
60	22.03	17.15		22.75	5.43	
65	22.33	17.37		23.19	5.38	
70	22.95	17.51		23.67		21.32
75	23.05	17.92		23.86		23.29
80	22.90	17.96		24.32		24.86
85	23.25	17.94		24.60		25.54
90	23.91	18.58		24.35		26.32
95	23.87	18.28		24.89		26.70
100	23.69	18.23		25.12		26.85
105	24.21	18.60		25.25		27.38
110	23.91	18.53		25.41		27.67
115	23.90		24.31	25.48		27.82
120	24.33		24.93	25.76		28.06
125	24.08		25.55	25.89		27.91
130	24.34		25.79	26.04		28.01
135	24.17		26.21	26.17		28.03
140	24.27		26.11	26.13		27.98
145	24.28		26.33	26.02		28.28
150	24.44		26.23	26.56		28.26
155	24.15		26.21	26.35		28.31
160	24.50		26.45	26.30		28.24
165	24.76		26.53	26.76		28.15
170	24.56		26.40	26.84		28.22
175	24.29		26.17	26.65	-	-
180	24.09		26.23	26.99	-	-
185	24.34		26.21	27.00	-	-
190	24.27		26.16	26.90	-	-
195	24.32		26.21	27.05	-	-
200	-	-	-	27.17	-	-
205	-	-	-	27.49	-	-
210	-	-	-	27.29	-	-
215	-	-	-	27.10	-	-
220	-	-	-	27.18	-	-
225	-	-	-	27.01	-	-
230	-	-	-	27.26	-	-
235	-	-	-	27.44	-	-
avarage	25.05	18.59	26.83	27.75	5.13	28.83

Table 10: Pre-training on the authentic data and fine-tuning on the synthetic data for Vietnamese-English NMT. BLEU Scores for each Checkpoint of the Models

training step (thousands)	<i>pre-train</i>	<i>fine-tune</i>
5	2.50	-
10	4.25	-
15	5.54	-
20	6.68	-
25	10.87	-
30	16.89	-
35	19.47	-
40	20.37	-
45	21.09	-
50	20.13	-
55	21.09	-
60	21.11	-
65	21.19	-
70	20.17	-
75	19.41	20.00
80	-	17.60
85	-	17.57
90	-	18.16
95	-	17.62
100	-	18.95
105	-	17.98
110	-	18.04
115	-	18.63
120	-	18.36
125	-	18.20
130	-	18.25
135	-	18.46
140	-	18.09
145	-	18.32
150	-	17.94
155	-	18.21
160	-	18.25
165	-	18.38
170	-	17.69
average	-	18.91

Table 11: Using different ratios of authentic to synthetic parallel data and its effect on the performance of Vietnamese-English NMT. Evaluation scores (BLEU) on the test set for each checkpoint (*tag-less_bt* colour code: BLACK – pre-train, RED – fine-tune)

training step (thousands)	<i>tagged_bt</i>			<i>tag-less_bt (joint BPE)</i>		
	1:1	1:3	1:5	1:1	1:3	1:5
5	2.31	1.78	1.87	1.65	2.87	1.82
10	3.31	3.64	3.23	8.68	9.78	6.16
15	7.35	9.88	4.05	11.93	13.06	12.81
20	14.45	15.29	5.08	13.52	14.11	14.20
25	17.12	16.99	5.71	14.36	15.03	13.63
30	18.54	19.22	13.36	15.24	15.92	16.08
35	20.45	19.62	16.41	15.49	15.95	15.24
40	21.76	19.71	17.28	15.90	16.39	16.34
45	21.86	20.99	19.20	16.18	17.08	17.32
50	21.96	22.03	19.93	16.05	17.15	16.36
55	22.87	21.85	20.33	16.52	17.29	17.45
60	22.71	22.03	20.85	16.58	17.15	16.99
65	23.05	22.33	21.62	16.30	17.37	17.83
70	23.33	22.95	21.83	16.82	17.51	16.94
75	22.84	23.05	21.84	16.80	17.92	18.36
80	23.25	22.90	23.05	16.76	17.96	17.86
85	23.51	23.25	22.56	16.54	17.94	18.14
90	22.99	23.91	23.15	16.47	18.58	18.44
95	23.06	23.87	23.62	16.69	18.28	18.52
100	22.89	23.69	23.37	23.04	18.23	18.31
105	22.96	24.21	23.58	24.05	18.60	18.80
110	22.80	23.91	23.49	24.81	18.53	18.42
115	22.92	23.90	23.90	24.87	24.31	19.22
120	22.83	24.33	23.69	25.50	24.93	19.07
125	22.29	24.08	23.79	24.86	25.55	18.47
130	22.53	24.34	24.10	26.29	25.79	19.44
135	22.43	24.17	24.01	25.77	26.21	19.48
140	22.36	24.27	24.52	25.95	26.11	19.16
145	22.34	24.28	24.57	25.94	26.33	19.14
150	22.08	24.44	23.88	25.96	26.23	19.01
155	22.08	24.15	24.62	25.76	26.21	24.82
160	21.74	24.50	24.49	-	26.45	25.79
165	22.14	24.76	24.12	-	26.53	26.21
170	22.73	24.56	25.47	-	26.40	26.27
175	-	24.29	-	-	26.17	26.51
180	-	24.09	-	-	26.23	26.78
185	-	24.34	-	-	26.21	26.91
190	-	24.27	-	-	26.16	26.60
195	-	24.32	-	-	26.21	26.57
200	-	-	-	-	-	26.58
205	-	-	-	-	-	26.64
210	-	-	-	-	-	25.70
215	-	-	-	-	-	26.29
average	22.73	25.05	25.47	26.15	26.83	27.14

Table 12: Fine-tuning the tagged and standard back-translations

training step (thousands)	<i>tagged.bt (1:5)</i>	<i>standard.bt (1:3)</i>
165	23.38	24.21
170	25.11	24.61
175	25.55	24.87
180	24.99	25.15
185	24.79	24.11
190	23.58	23.98
195	23.37	23.36
200	23.15	22.97
average	25.64	25.32

Table 13: This table shows how often a conclusion with 95% statistical significance is made for comparing the various approaches. We used different sample sizes of 100, 500 and 1000 sentences for each of the approach on English-Vietnamese and English-German low resource NMT.

System Comparison	Vi-En				De-En			
	BLEU difference	Sample size			BLEU difference	Sample size		
		100	500	1000		100	500	1000
Standard BT is better than baseline	3.06	100%	100%	100%	4.92	100%	100%	100%
Tagged BT is better than baseline	2.83	100%	100%	100%	6.81	100%	100%	100%
Tagged BT is better than Standard BT	-0.23	27%	28%	30.2%	1.89	100%	100%	100%
Tag-less BT is better than baseline	3.55	100%	100%	100%	5.09	100%	100%	100%
Tag-less BT better than Standard BT	0.49	92.8%	94.6%	94.8%	0.16	87.1%	87.4%	84%
Tag-less BT is better than Tagged BT	0.72	98.8%	98.2%	98%	-1.72	0%	0%	0%
Tag-less BT (joint BPE) is better than baseline	4.61	100%	100%	100%	7.89	100%	100%	100%
Tag-less BT (joint BPE) is better than Standard BT	1.55	100%	100%	100%	2.96	100%	100%	100%
Tag-less BT (joint BPE) is better than Tagged BT	1.78	100%	100%	100%	1.08	100%	100%	100%

Figures

Figure 1

Tag-less Back-Translation: Training the forward model on the synthetic parallel data generated using the backward model. The forward model is then ne-tuned on the authentic data.

Figure 2

Tag-less back-translation: pre-training on synthetic data and ne-tuning on authentic data. Showing how this technique compares to the baseline and the standard back-translation approaches on the test set.

Figure 3

Tagged Vs Tag-less back-translation.

Figure 4

Fine-tuning the baseline model on the synthetic data. Evaluation scores on the test set.