

# When and Why is Unsupervised Neural Machine Translation Useless?

Yunsu Kim Miguel Graça<sup>†</sup> Hermann Ney

Human Language Technology and Pattern Recognition Group

RWTH Aachen University, Aachen, Germany

{surname}@cs.rwth-aachen.de

## Abstract

This paper studies the practicality of the current state-of-the-art unsupervised methods in neural machine translation (NMT). In ten translation tasks with various data settings, we analyze the conditions under which the unsupervised methods fail to produce reasonable translations. We show that their performance is severely affected by linguistic dissimilarity and domain mismatch between source and target monolingual data. Such conditions are common for low-resource language pairs, where unsupervised learning works poorly. In all of our experiments, supervised and semi-supervised baselines with 50k-sentence bilingual data outperform the best unsupervised results. Our analyses pinpoint the limits of the current unsupervised NMT and also suggest immediate research directions.

## 1 Introduction

Statistical methods for machine translation (MT) require a large set of sentence pairs in two languages to build a decent translation system (Resnik and Smith, 2003; Koehn, 2005). Such bilingual data is scarce for most language pairs and its quality varies largely over different domains (Al-Onaizan et al., 2002; Chu and Wang, 2018). Neural machine translation (NMT) (Bahdanau et al., 2015; Vaswani et al., 2017), the standard paradigm of MT these days, has been claimed to suffer from the data scarcity more severely than phrase-based MT (Koehn and Knowles, 2017).

Unsupervised NMT, which trains a neural translation model only with monolingual corpora, was

proposed for those scenarios which lack bilingual data (Artetxe et al., 2018b; Lample et al., 2018a). Despite its progress in research, the performance of the unsupervised methods has been evaluated mostly on high-resource language pairs, e.g. German $\leftrightarrow$ English or French $\leftrightarrow$ English (Artetxe et al., 2018b; Lample et al., 2018a; Yang et al., 2018; Artetxe et al., 2018a; Lample et al., 2018b; Ren et al., 2019b; Artetxe et al., 2019; Sun et al., 2019; Sen et al., 2019). For these language pairs, huge bilingual corpora are already available, so there is no need for unsupervised learning in practice. Empirical results in these tasks do not carry over to low-resource language pairs; they simply fail to produce any meaningful translations (Neubig and Hu, 2018; Guzmán et al., 2019).

This paper aims for a more comprehensive and pragmatic study on the performance of unsupervised NMT. Our experiments span ten translation tasks in the following five language pairs:

- German $\leftrightarrow$ English: similar languages, abundant bilingual/monolingual data
- Russian $\leftrightarrow$ English: distant languages, abundant bilingual/monolingual data, similar sizes of the alphabet
- Chinese $\leftrightarrow$ English: distant languages, abundant bilingual/monolingual data, very different sizes of the alphabet
- Kazakh $\leftrightarrow$ English: distant languages, scarce bilingual data, abundant monolingual data
- Gujarati $\leftrightarrow$ English: distant languages, scarce bilingual/monolingual data

For each task, we compare the unsupervised performance with its supervised and semi-supervised counterparts. In addition, we make the monolingual training data vary in size and domain to cover many more scenarios, showing under which conditions unsupervised NMT works poorly.

Here is a summary of our contributions:

<sup>†</sup> The author is now at DeepL GmbH.

© 2020 The authors. This article is licensed under a Creative Commons 3.0 licence, no derivative works, attribution, CC-BY-ND.

- We thoroughly evaluate the performance of state-of-the-art unsupervised NMT in numerous real and artificial translation tasks.
- We provide guidelines on whether to employ unsupervised NMT in practice, by showing how much bilingual data is sufficient to outperform the unsupervised results.
- We clarify which factors make unsupervised NMT weak and which points must be improved, by analyzing the results both quantitatively and qualitatively.

## 2 Related Work

The idea of unsupervised MT dates back to word-based decipherment methods (Knight et al., 2006; Ravi and Knight, 2011). They learn only lexicon models at first, but add alignment models (Dou et al., 2014; Nuhn, 2019) or heuristic features (Naim et al., 2018) later. Finally, Artetxe et al. (2018a) and Lample et al. (2018b) train a fully-fledged phrase-based MT system in an unsupervised way.

With neural networks, unsupervised learning of a sequence-to-sequence NMT model has been proposed by Lample et al. (2018a) and Artetxe et al. (2018b). Though having slight variations (Yang et al., 2018; Sun et al., 2019; Sen et al., 2019), unsupervised NMT approaches commonly 1) learn a shared model for both source→target and target→source 2) using iterative back-translation, along with 3) a denoising autoencoder objective. They are initialized with either cross-lingual word embeddings or a cross-lingual language model (LM). To further improve the performance at the cost of efficiency, Lample et al. (2018b), Ren et al. (2019b) and Artetxe et al. (2019) combine unsupervised NMT with unsupervised phrase-based MT. On the other hand, one can also avoid the long iterative training by applying a separate denoiser directly to the word-by-word translations from cross-lingual word embeddings (Kim et al., 2018; Pourdamghani et al., 2019).

Unsupervised NMT approaches have been so far evaluated mostly on high-resource language pairs, e.g. French→English, for academic purposes. In terms of practicality, they tend to underperform in low-resource language pairs, e.g. Azerbaijani→English (Neubig and Hu, 2018) or Nepali→English (Guzmán et al., 2019). To the best of our knowledge, this work is the first to systematically evaluate and analyze unsupervised learning for NMT in various data settings.

## 3 Unsupervised NMT

This section reviews the core concepts of the recent unsupervised NMT framework and describes to which points they are potentially vulnerable.

### 3.1 Bidirectional Modeling

Most of the unsupervised NMT methods share the model parameters between source→target and target→source directions. They also often share a joint subword vocabulary across the two languages (Sennrich et al., 2016b).

Sharing a model among different translation tasks has been shown to be effective in multilingual NMT (Firat et al., 2016; Johnson et al., 2017; Aharoni et al., 2019), especially in improving performance on low-resource language pairs. This is due to the commonality of natural languages; learning to represent a language is helpful to represent other languages, e.g. by transferring knowledge of general sentence structures. It also provides good regularization for the model.

Unsupervised learning is an extreme scenario of MT, where bilingual information is very weak. To supplement the weak and noisy training signal, knowledge transfer and regularization are crucial, which can be achieved by the bidirectional sharing. It is based on the fact that a translation problem is dual in nature; source→target and target→source tasks are conceptually related to each other.

Previous works on unsupervised NMT vary in the degree of sharing: the whole encoder (Artetxe et al., 2018b; Sen et al., 2019), the middle layers (Yang et al., 2018; Sun et al., 2019), or the whole model (Lample et al., 2018a; Lample et al., 2018b; Ren et al., 2019a; Conneau and Lample, 2019).

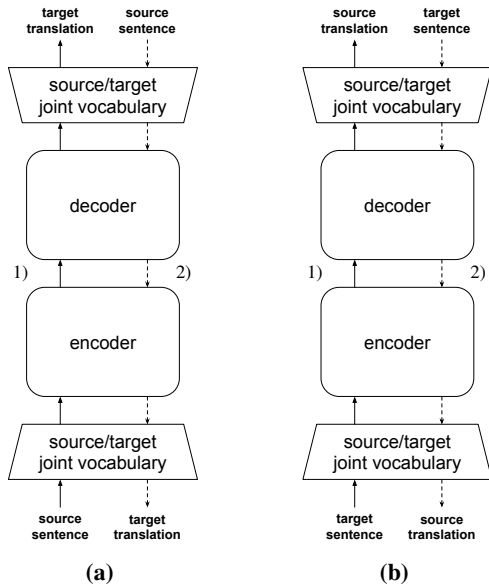
Note that the network sharing is less effective among linguistically distinct languages in NMT (Kocmi and Bojar, 2018; Kim et al., 2019a). It still works as a regularizer, but transferring knowledge is harder if the morphology or word order is quite different. We show how well unsupervised NMT performs on such language pairs in Section 4.1.

### 3.2 Iterative Back-Translation

Unsupervised learning for MT assumes no bilingual data for training. A traditional remedy for the data scarcity is generating synthetic bilingual data from monolingual text (Koehn, 2005; Schwenk, 2008; Sennrich et al., 2016a). To train a bidirectional model of Section 3.1, we need bilingual data of both translation directions. Therefore, most un-

supervised NMT methods back-translate in both directions, i.e. source and target monolingual data to target and source language, respectively.

In unsupervised learning, the synthetic data should be created not only once at the beginning but also repeatedly throughout the training. At the early stages of training, the model might be too weak to generate good translations. Hence, most methods update the training data as the model gets improved during training. The improved model for source→target direction back-translates source monolingual data, which improves the model for target→source direction, and vice versa. This cycle is called dual learning (He et al., 2016) or iterative back-translation (Hoang et al., 2018). Figure 1 shows the case when it is applied to a fully shared bidirectional model.



**Figure 1:** Iterative back-translation for training a bidirectional sequence-to-sequence model. The model first translates monolingual sentences (solid arrows), and then gets trained with the translation as the input and the original as the output (dashed arrows). This procedure alternates between (a) source→target and (b) target→source translations.

One can tune the amount of back-translations per iteration: a mini-batch (Artetxe et al., 2018b; Yang et al., 2018; Conneau and Lample, 2019; Ren et al., 2019a), the whole monolingual data (Lample et al., 2018a; Lample et al., 2018b; Sun et al., 2019), or some size in between (Artetxe et al., 2019; Ren et al., 2019b).

However, even if carefully scheduled, the iterative training cannot recover from a bad optimum if the initial model is too poor. Experiments in Section 4.5 highlight such cases.

### 3.3 Initialization

To kickstart the iterative training, the model should be able to generate meaningful translations already in the first iteration. We cannot expect the training to progress from a randomly initialized network and the synthetic data generated by it.

Cross-lingual embeddings give a good starting point for the model by defining a joint continuous space shared by multiple languages. Ideally, in such a space, close embedding vectors are semantically related to each other regardless of their languages; they can be possible candidates for translation pairs (Mikolov et al., 2013). It can be learned either in word level (Artetxe et al., 2017; Conneau et al., 2018) or in sentence level (Conneau and Lample, 2019) using only monolingual corpora.

In the word level, we can initialize the embedding layers with cross-lingual word embedding vectors (Artetxe et al., 2018b; Lample et al., 2018a; Yang et al., 2018; Lample et al., 2018b; Artetxe et al., 2019; Sun et al., 2019). On the other hand, the whole encoder/decoder parameters can be initialized with cross-lingual sequence training (Conneau and Lample, 2019; Ren et al., 2019a; Song et al., 2019).

Cross-lingual word embedding has limited performance among distant languages (Søgaard et al., 2018; Nakashole and Flauger, 2018) and so does cross-lingual LM (Pires et al., 2019). Section 4.5 shows the impact of a poor initialization.

### 3.4 Denoising Autoencoder

Initializing the word embedding layers furnishes the model with cross-lingual matching in the lexical embedding space, but does not provide any information on word orders or generation of text. Cross-lingual LMs encode word sequences in different languages, but they are not explicitly trained to reorder source words to the target language syntax. Both ways do not initialize the crucial parameters for reordering: the encoder-decoder attention and the recurrence on decoder states.


As a result, an initial model for unsupervised NMT tends to generate word-by-word translations with little reordering, which are very non-fluent when source and target languages have distinct word orders. Training on such data discourages the model from reordering words, which might cause a vicious cycle by generating even less-reordered synthetic sentence pairs in the next iterations.

Accordingly, unsupervised NMT employs an

		de-en		ru-en		zh-en		kk-en		gu-en	
		German	English	Russian	English	Chinese	English	Kazakh	English	Gujarati	English
Language family		Germanic	Germanic	Slavic	Germanic	Sinitic	Germanic	Turkic	Germanic	Indic	Germanic
Alphabet Size		60	52	66	52	8,105	52	42	52	91	52
Monolingual	Sentences	100M		71.6M		30.8M		18.5M		4.1M	
	Words	1.8B	2.3B	1.1B	2.0B	1.4B	699M	278.5M	421.5M	121.5M	93.8M
Bilingual	Sentences	5.9M		25.4M		18.9M		222k		156k	
	Words	137.4M	144.9M	618.6M	790M	440.3M	482.9M	1.6M	1.9M	2.3M	1.5M

**Table 1:** Training data statistics.

additional training objective of denoising autoencoding (Hill et al., 2016). Given a clean sentence, artificial noises are injected, e.g. deletion or permutation of words, to make a corrupted input. The denoising objective trains the model to reorder the noisy input to the correct syntax, which is essential for generating fluent outputs. This is done for each language individually with monolingual data, as shown in Figure 2.



**Figure 2:** Denoising autoencoder training for source or target language.

Once the model is sufficiently trained for denoising, it is helpful to remove the objective or reduce its weight (Graça et al., 2018). At the later stages of training, the model gets improved in re-ordering and translates better; learning to denoise might hurt the performance in clean test sets.

## 4 Experiments and Analysis

**Data** Our experiments were conducted on WMT 2018 German↔English and Russian↔English, WMT 2019 Chinese↔English, Kazakh↔English, and Gujarati↔English (Table 1). We pre-

processed the data using the MOSES<sup>1</sup> tokenizer and a frequent caser. For Chinese, we used the JIEBA segmenter<sup>2</sup>. Lastly, byte pair encoding (BPE) (Sennrich et al., 2016b) was learned jointly over source and target languages with 32k merges and applied without vocabulary threshold.

**Model** We used 6-layer Transformer base architecture (Vaswani et al., 2017) by default: 512-dimension embedding/hidden layers, 2048-dimension feedforward sublayers, and 8 heads.

**Decoding and Evaluation** Decoding was done with beam size 5. We evaluated the test performance with SACREBLEU (Post, 2018).

**Unsupervised Learning** We ran XLM<sup>3</sup> by Conneau and Lample (2019) for the unsupervised experiments. The back-translations were done with beam search for each mini-batch of 16k tokens. The weight of the denoising objective started with 1 and linearly decreased to 0.1 until 100k updates, and then decreased to 0 until 300k updates.

The model’s encoder and decoder were both initialized with the same pre-trained cross-lingual LM. We removed the language embeddings from the encoder for better cross-linguality (see Section 4.6). Unless otherwise specified, we used the same monolingual training data for both pre-training and translation training. For the pre-training, we set the batch size to 256 sentences (around 66k tokens).

Training was done with Adam (Kingma and Ba, 2014) with an initial learning rate of 0.0001, where dropout (Srivastava et al., 2014) of probability 0.1 was applied to each layer output and attention components. With a checkpoint frequency of 200k sentences, we stopped the training when the validation perplexity (pre-training) or BLEU (translation training) was not improved for ten check-

<sup>1</sup><http://www.statmt.org/moses>

<sup>2</sup><https://github.com/fxsjy/jieba>

<sup>3</sup><https://github.com/facebookresearch/XLM>

Approach	BLEU [%]									
	de-en	en-de	ru-en	en-ru	zh-en	en-zh	kk-en	en-kk	gu-en	en-gu
Supervised	39.5	39.1	29.1	24.7	26.2	39.6	10.3	2.4	9.9	3.5
Semi-supervised	43.6	41.0	30.8	28.8	25.9	42.7	12.5	3.1	14.2	4.0
Unsupervised	23.8	20.2	12.0	9.4	1.5	2.5	2.0	0.8	0.6	0.6

**Table 2:** Comparison among supervised, semi-supervised, and unsupervised learning. All bilingual data was used for the (semi-)supervised results and all monolingual data was used for the unsupervised results (see Table 1). All results are computed on newstest2019 of each task, except for de-en/en-de and ru-en/en-ru on newstest2018.

points. We extensively tuned the hyperparameters for a single GPU with 12GB memory, which is widely applicable to moderate industrial/academic environments. All other hyperparameter values follow the recommended settings of XLM.

**Supervised Learning** Supervised experiments used the same hyperparameters as the unsupervised learning, except 12k tokens for the batch size, 0.0002 for the initial learning rate, and 10k batches for each checkpoint.


If the bilingual training data contains less than 500k sentence pairs, we reduced the BPE merges to 8k, the batch size to 2k, and the checkpoint frequency to 4k batches; we also increased the dropout rate to 0.3 (Sennrich and Zhang, 2019).

**Semi-supervised Learning** Semi-supervised experiments continued the training from the supervised baseline with back-translations added to the training data. We used 4M back-translated sentences for the low-resource cases, i.e. if the original bilingual data has less than 500k lines, and 10M back-translated sentences otherwise.

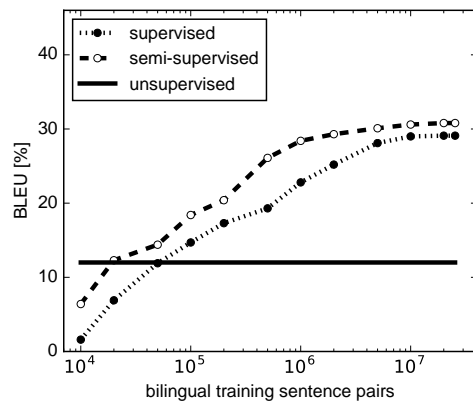
#### 4.1 Unsupervised vs. (Semi-)Supervised

We first address the most general question of this paper: For NMT, can unsupervised learning replace semi-supervised or supervised learning? Table 2 compares the unsupervised performance to simple supervised and semi-supervised baselines.

In all tasks, unsupervised learning shows much worse performance than (semi-)supervised learning. It produces readable translations in two high-resource language pairs (German↔English and Russian↔English), but their scores are only around half of the semi-supervised systems. In other three language pairs, unsupervised NMT fails to converge at any meaningful optimum, reaching less than 3% BLEU scores. Note that, in these three tasks, source and target languages are very different in the alphabet, morphology, and



(a) German→English



(b) Russian→English

**Figure 3:** Supervised and semi-supervised learning over bilingual training data size. Unsupervised learning (horizontal line) uses all monolingual data of Table 1.

word order, etc. The results in Kazakh↔English and Gujarati↔English show that the current unsupervised NMT cannot be an alternative to (semi-)supervised NMT in low-resource conditions.

To discover the precise condition where the unsupervised learning is useful in practice, we vary the size of the given bilingual training data for (semi-)supervised learning and plot the results in Figure 3. Once we have 50k bilingual sentence pairs in German↔English, simple semi-supervised learning already outperforms unsupervised learning with 100M monolingual sentences

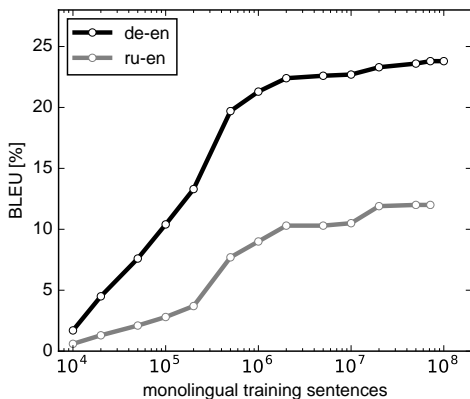
in each language. Even without back-translations (supervised), 100k-sentence bilingual data is sufficient to surpass unsupervised NMT.

In the Russian↔English task, the unsupervised learning performance can be more easily achieved with only 20k bilingual sentence pairs using semi-supervised learning. This might be due to that Russian and English are more distant to each other than German and English, thus bilingual training signal is more crucial for Russian↔English.

Note that for these two language pairs, the bilingual data for supervised learning are from many different text domains, whereas the monolingual data are from exactly the same domain of the test sets. Even with such an advantage, the large-scale unsupervised NMT cannot compete with supervised NMT with tiny out-of-domain bilingual data.

## 4.2 Monolingual Data Size

In this section, we analyze how much monolingual data is necessary to make unsupervised NMT produce reasonable performance. Figure 4 shows the unsupervised results with different amounts of monolingual training data. We keep the equal size for source and target data, and the domain is also the same for both (web-crawled news).



**Figure 4:** Unsupervised NMT performance over the size of monolingual training data, where source and target sides have the same size.

For German→English, training with only 1M sentences already gives a reasonable performance, which is only around 2% BLEU behind the 100M-sentence case. The performance starts to saturate already after 5M sentences, with only marginal improvements by using more than 20M sentences. We observe a similar trend in Russian→English.

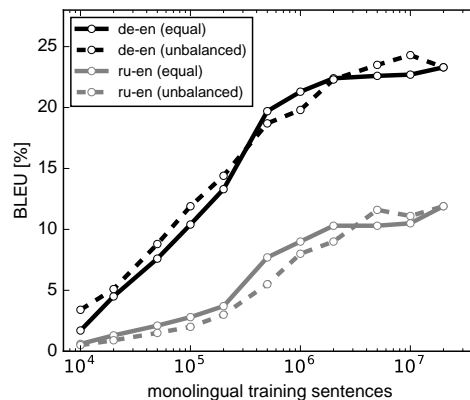
This shows that, for the performance of unsupervised NMT, using a massive amount of monolingual data is not as important as the similarity

of source and target languages. Comparing to supervised learning (see Figure 3), the performance saturates faster when increasing the training data, given the same model size.

## 4.3 Unbalanced Data Size

What if the size of available monolingual data is largely different for source and target languages? This is often the case for low-resource language pairs involving English, where there is plenty of data for English but not for the other side.

Our experiments so far intentionally use the same number of sentences for both sides. In Figure 5, we reduced the source data gradually while keeping the large target data fixed. To counteract the data imbalance, we oversampled the smaller side to make the ratio of source-target 1:1 for BPE learning and mini-batch construction (Conneau and Lample, 2019). We compare such unbalanced data settings to the previous equal-sized source/target settings.



**Figure 5:** Unsupervised NMT performance over source training data size, where the target training data is fixed to 20M sentences (dashed line). Solid line is the case where the target data has the same number of sentences as the source side.

Interestingly, when we decrease the target data accordingly (balanced, solid line), the performance is similar or sometimes better than using the full target data (unbalanced, dashed line). This means that it is not beneficial to use oversized data on one side in unsupervised NMT training.

If the data is severely unbalanced, the distribution of the smaller side should be much sparser than that of the larger side. The network tries to generalize more on the smaller data, reserving the model capacity for smoothing (Olson et al., 2018). Thus it learns to represent a very different distribution of each side, which is challenging in a shared model (Section 3.1). This could be the reason for

no merit in using larger data on one side.

#### 4.4 Domain Similarity

In high-resource language pairs, it is feasible to collect monolingual data of the same domain on both source and target languages. However, for low-resource language pairs, it is difficult to match the data domain of both sides on a large scale. For example, our monolingual data for Kazakh is mostly from Wikipedia and Common Crawl, while the English data is solely from News Crawl. In this section, we study how the domain similarity of monolingual data on the two sides affects the performance of unsupervised NMT.

In Table 3, we artificially change the domain of the source side to politics (UN Corpus<sup>4</sup>) or random (Common Crawl), while keeping the target domain fixed to newswire (News Crawl). The results show that the domain matching is critical for unsupervised NMT. For instance, although German and English are very similar languages, we see the performance of German↔English deteriorate down to -11.8% BLEU by the domain mismatch.

Domain (en)	Domain (de/ru)	BLEU [%]			
		de-en	en-de	ru-en	en-ru
	Newswire	23.3	19.9	11.9	9.3
Newswire	Politics	11.5	12.2	2.3	2.5
	Random	18.4	16.4	6.9	6.1

**Table 3:** Unsupervised NMT performance where source and target training data are from different domains. The data size on both sides is the same (20M sentences).

Table 4 shows a more delicate case where we keep the same domain for both sides (newswire) but change the providers and years of the news articles. Our monolingual data for Chinese (Table 1) consist mainly of News Crawl (from years 2008-2018) and Gigaword 4th edition (from years 1995-2008). We split out the News Crawl part (1.7M sentences) and trained an unsupervised NMT model with the same amount of English monolingual data (from News Crawl 2014-2017). Surprisingly, this experiment yields much better results than using all available data. Even if the size is small, the source and target data are collected in the same way (web-crawling) from similar years (2010s), which seems to be crucial for unsupervised NMT to work.

On the other hand, when using the Gigaword part (28.6M sentences) on Chinese, unsupervised

<sup>4</sup><https://conferences.unite.un.org/uncorpus>

Years (en)	Years (zh)	#sents (en/zh)	BLEU [%]	
			zh-en	en-zh
2014-2017	2008-2018	1.7M	5.4	15.1
	1995-2008	28.6M	1.5	1.9

**Table 4:** Unsupervised NMT performance where source and target training data are from the same domain (newswire) but different years.

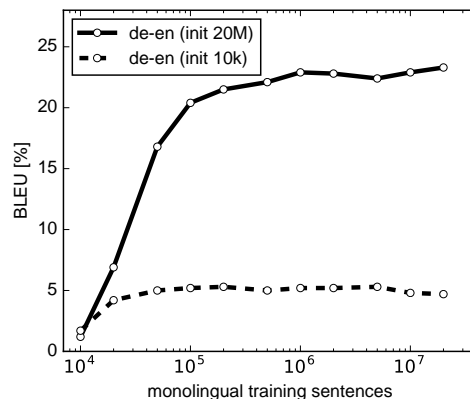
learning again does not function properly. Now the source and target text are from different decades; the distribution of topics might be different. Also, the Gigaword corpus is from traditional newspaper agencies which can have a different tone from the online text of News Crawl. Despite the large scale, unsupervised NMT proves to be sensitive to a subtle discrepancy of topic, style, period, etc. between source and target data.

These results agree with Søgaard et al. (2018) who show that modern cross-lingual word embedding methods fail in domain mismatch scenarios.

#### 4.5 Initialization vs. Translation Training

Thus far, we have seen a number of cases where unsupervised NMT breaks down. But which part of the learning algorithm is more responsible for the performance: initialization (Section 3.3) or translation training (Section 3.2 and 3.4)?

In Figure 6, we control the level of each of the two training stages and analyze its impact on the final performance. We pre-trained two cross-lingual LMs as initializations of different quality: bad (using 10k sentences) and good (using 20M sentences). For each initial point, we continued the translation training with different amounts of data from 10k to 20M sentences.



**Figure 6:** Unsupervised NMT performance over the training data size for translation training, where the pre-training data for initialization is fixed (10k or 20M sentences).

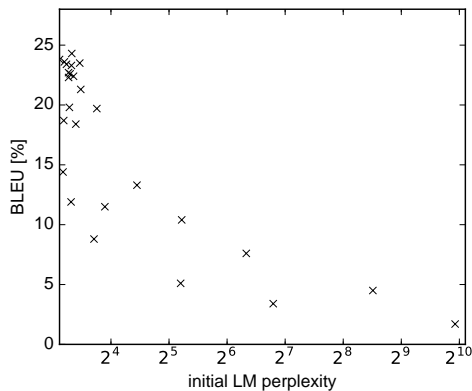
From the bad initialization, unsupervised learning cannot build a reasonable NMT model, no mat-

Task	BLEU [%]	Source input	System output	Reference output
de-en	23.8	Seit der ersten <u>Besichtigung</u> wurde die 1.000 <u>Quadratfuß</u> große ...	Since the first <u>Besichtigung</u> , the 3,000 <u>square</u> fueled ...	Since the first viewing, the 1,000sq ft flat has ...
	10.4	<u>München</u> 1856: <u>Vier</u> Karten, die Ihren Blick auf die <u>Stadt</u> verändern	<u>Australia</u> 1856: <u>Eight</u> things that can keep your way to the <u>UK</u>	Munich 1856: Four maps that will change your view of the city
ru-en	12.0	В ходе первоочередных оператив- но-следственных мероприятий ус- тановлена личность роженицы	The первоочередных оператив- но-следственных мероприятий have been established by the dolphin	The identity of the mother was de- termined during preliminary inves- tigative and operational measures
zh-en	1.5	... 调整要兼顾生产需要和消费需求。	... 调整要兼顾生产需要 and 消费需 求。	... adjustment must balance produc- tion needs with consumer demands.

**Table 5:** Problematic translation outputs from unsupervised NMT systems (input copying, ambiguity in the same context).

ter how much data is used in translation training. When the initial model is strong, it is possible to reach 20% BLEU by translation training with only 100k sentences. Using 1M sentences in translation training, the performance is already comparable to its best. Once the model is pre-trained well for cross-lingual representations, fine-tuning the translation-specific components seems manageable with relatively small data.

This demonstrates the importance of initialization over translation training in the current unsupervised NMT. Translation training relies solely on model-generated inputs, i.e. back-translations, which do not reflect the true distribution of the input language when generated with a poor initial model. On Figure 7, we plot all German→English unsupervised results we conducted up to the previous section. It shows that the final performance generally correlates with the initialization quality.



**Figure 7:** Unsupervised NMT performance over the validation perplexity of the initial cross-lingual LM (de-en).

#### 4.6 Qualitative Examples

In this section, we analyze translation outputs of unsupervised systems to find out why they record such low BLEU scores. Do unsupervised systems have particular problems in the outputs other than limited adequacy/fluency?

Table 5 shows translation examples from the unsupervised systems. The first notable problem is copying input words to the output. This happens when the encoder has poor cross-linguality, i.e. does not concurrently model two languages well in a shared space. The decoder then can easily detect the input language by reading the encoder and may emit output words in the same language.

A good cross-lingual encoder should not give away information on the input language to the decoder. The decoder must instead rely on the output language embeddings or an indicator token (e.g.  $\langle 2_{en} \rangle$ ) to determine the language of output tokens. As a simple remedy, we removed the language embeddings from the encoder and obtained consistent improvements, e.g. from 4.3% to 11.9% BLEU in Russian→English. However, the problem still remains partly even in our best-performing unsupervised system (the first example).

The copying occurs more often in inferior systems (the last example), where the poor initial cross-lingual LM is the main reason for the worse performance (Section 4.5). Note that the auto-encoding (Section 3.4) also encourages the model to generate outputs in the input language.

Another problem is that the model cannot distinguish words that appear in the same context. In the second example, the model knows that *Vier* in German (*Four* in English) is a number, but it generates a wrong number in English (*Eight*). The initial LM is trained to predict either *Four* or *Eight* given the same surrounding words (e.g. 1856, things) and has no clue to map *Four* to *Vier*.

The model cannot learn these mappings by itself with back-translations. This problem can be partly solved by subword modeling (Bojanowski et al., 2017) or orthographic features (Riley and Gildea, 2018; Artetxe et al., 2019), which are however not effective for language pairs with disjoint alphabets.



## 5 Conclusion and Outlook

In this paper, we examine the state-of-the-art unsupervised NMT in a wide range of tasks and data settings. We find that the performance of unsupervised NMT is seriously affected by these factors:

- Linguistic similarity of source and target languages
- Domain similarity of training data between source and target languages

It is very hard to fulfill these in low-/zero-resource language pairs, which makes the current unsupervised NMT useless in practice. We also find that the performance is not improved by using massive monolingual data on one or both sides.

In practice, a simple, non-tuned semi-supervised baseline with only less than 50k bilingual sentence pairs is sufficient to outperform our best large-scale unsupervised system. At this moment, we cannot recommend unsupervised learning for building MT products if there are at least small bilingual data.

For the cases where there is no bilingual data available at all, we plan to systematically compare the unsupervised NMT to pivot-based methods (Kim et al., 2019b; Currey and Heafield, 2019) or multilingual zero-shot translation (Johnson et al., 2017; Aharoni et al., 2019).

To make unsupervised NMT useful in the future, we suggest the following research directions:

**Language-/Domain-agnostic LM** We show in Section 4.5 that the initial cross-lingual LM actually determines the performance of unsupervised NMT. In Section 4.6, we argue that the poor performance is due to input copying, for which we blame a poor cross-lingual LM. The LM pre-training must therefore handle dissimilar languages and domains equally well. This might be done by careful data selection or better regularization methods.

**Robust Translation Training** On the other hand, the current unsupervised NMT lacks a mechanism to bootstrap out of a poor initialization. Inspired by classical decipherment methods (Section 2), we might devalue noisy training examples or artificially simplify the problem first.

## References

Aharoni, Roei, Melvin Johnson, and Orhan Firat. 2019. Massively multilingual neural machine translation. In *NAACL-HLT*, pages 3874–3884.

- Al-Onaizan, Yaser, Ulrich Germann, Ulf Hermjakob, Kevin Knight, Philipp Koehn, Daniel Marcu, and Kenji Yamada. 2002. Translation with scarce bilingual resources. *Machine Translation*, 17(1):1–17.
- Artetxe, Mikel, Gorka Labaka, and Eneko Agirre. 2017. Learning bilingual word embeddings with (almost) no bilingual data. In *ACL*, pages 451–462.
- Artetxe, Mikel, Gorka Labaka, and Eneko Agirre. 2018a. Unsupervised statistical machine translation. In *EMNLP*, page 3632–3642.
- Artetxe, Mikel, Gorka Labaka, Eneko Agirre, and Kyunghyun Cho. 2018b. Unsupervised neural machine translation. In *ICLR*.
- Artetxe, Mikel, Gorka Labaka, and Eneko Agirre. 2019. An effective approach to unsupervised machine translation. In *ACL*, pages 194–203.
- Bahdanau, Dzmitry, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Bojanowski, Piotr, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *TACL*, 5:135–146.
- Chu, Chenhui and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. In *COLING*, pages 1304–1319.
- Conneau, Alexis and Guillaume Lample. 2019. Cross-lingual language model pretraining. In *NeurIPS*, pages 7057–7067.
- Conneau, Alexis, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2018. Word translation without parallel data. In *ICLR*.
- Currey, Anna and Kenneth Heafield. 2019. Zero-resource neural machine translation with monolingual pivot data. In *WNGT*, pages 99–107.
- Dou, Qing, Ashish Vaswani, and Kevin Knight. 2014. Beyond parallel data: Joint word alignment and decipherment improves machine translation. In *EMNLP*, pages 557–565.
- Firat, Orhan, Kyunghyun Cho, and Yoshua Bengio. 2016. Multi-way, multilingual neural machine translation with a shared attention mechanism. In *NAACL-HLT*, pages 866–875.
- Graça, Miguel, Yunsu Kim, Julian Schamper, Jiahui Geng, and Hermann Ney. 2018. The RWTH aachen university English-German and German-English unsupervised neural machine translation systems for WMT 2018. In *WMT*.
- Guzmán, Francisco, Peng-Jen Chen, Myle Ott, Juan Pino, Guillaume Lample, Philipp Koehn, Vishrav Chaudhary, and Marc’Aurelio Ranzato. 2019. The FLORES evaluation datasets for low-resource machine translation: Nepali–English and Sinhala–English. In *EMNLP-IJCNLP*, pages 6097–6110.
- He, Di, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. Dual learning for machine translation. In *NIPS*, pages 820–828.
- Hill, Felix, Kyunghyun Cho, and Anna Korhonen. 2016. Learning distributed representations of sentences from unlabelled data. In *NAACL-HLT*, pages 1367–1377.
- Hoang, Vu Cong Duy, Philipp Koehn, Gholamreza Haffari, and Trevor Cohn. 2018. Iterative back-translation for neural machine translation. In *WNGT*, pages 18–24.

- Johnson, Melvin, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *TACL*, 5(1):339–351.
- Kim, Yunsu, Jiahui Geng, and Hermann Ney. 2018. Improving unsupervised word-by-word translation with language model and denoising autoencoder. In *EMNLP*, pages 862–868.
- Kim, Yunsu, Yingbo Gao, and Hermann Ney. 2019a. Effective cross-lingual transfer of neural machine translation models without shared vocabularies. In *ACL*, pages 1246–1257.
- Kim, Yunsu, Petre Petrov, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2019b. Pivot-based transfer learning for neural machine translation between non-English languages. In *EMNLP-IJCNLP*, pages 866–876.
- Kingma, Diederik P and Jimmy Ba. 2014. Adam: A method for stochastic optimization.
- Knight, Kevin, Anish Nair, Nishit Rathod, and Kenji Yamada. 2006. Unsupervised analysis for decipherment problems. In *COLING/ACL*, pages 499–506.
- Kocmi, Tom and Ondřej Bojar. 2018. Trivial transfer learning for low-resource neural machine translation. In *WMT*, pages 244–252.
- Koehn, Philipp and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *WNMT*, pages 28–39.
- Koehn, Philipp. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT Summit*, pages 79–86.
- Lample, Guillaume, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018a. Unsupervised machine translation using monolingual corpora only. In *ICLR*.
- Lample, Guillaume, Myle Ott, Alexis Conneau, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018b. Phrase-based & neural unsupervised machine translation. In *EMNLP*, pages 5039–5049.
- Mikolov, Tomas, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient estimation of word representations in vector space.
- Naim, Iftekhar, Parker Riley, and Daniel Gildea. 2018. Feature-based decipherment for machine translation. *Computational Linguistics*, 44(3):525–546.
- Nakashole, Ndapandula and Raphael Flauger. 2018. Characterizing departures from linearity in word translation. In *ACL*, pages 221–227.
- Neubig, Graham and Junjie Hu. 2018. Rapid adaptation of neural machine translation to new languages. In *EMNLP*, pages 875–880.
- Nuhn, Malte. 2019. *Unsupervised Training with Applications in Natural Language Processing*. Ph.D. thesis, Computer Science Department, RWTH Aachen University.
- Olson, Matthew, Abraham Wyner, and Richard Berk. 2018. Modern neural networks generalize on small data sets. In *NIPS*, pages 3619–3628.
- Pires, Telmo, Eva Schlinger, and Dan Garrette. 2019. How multilingual is multilingual bert? In *ACL*, pages 4996–5001.
- Post, Matt. 2018. A call for clarity in reporting bleu scores. In *WMT*, pages 186–191.
- Pourdamghani, Nima, Nada Aldarrab, Marjan Ghazvininejad, Kevin Knight, and Jonathan May. 2019. Translating translationese: A two-step approach to unsupervised machine translation. In *ACL*, pages 3057–3062.
- Ravi, Sujith and Kevin Knight. 2011. Deciphering foreign language. In *ACL*, pages 12–21.
- Ren, Shuo, Yu Wu, Shujie Liu, Ming Zhou, and Shuai Ma. 2019a. Explicit cross-lingual pre-training for unsupervised machine translation. In *EMNLP-IJCNLP*, pages 770–779.
- Ren, Shuo, Zhirui Zhang, Shujie Liu, Ming Zhou, and Shuai Ma. 2019b. Unsupervised neural machine translation with smt as posterior regularization.
- Resnik, Philip and Noah A. Smith. 2003. The web as a parallel corpus. *Computational Linguistics*, 29(3):349–380.
- Riley, Parker and Daniel Gildea. 2018. Orthographic features for bilingual lexicon induction. In *ACL*, pages 390–394.
- Schwenk, Holger. 2008. Investigations on large-scale lightly-supervised training for statistical machine translation. In *IWSLT*.
- Sen, Sukanta, Kamal Kumar Gupta, Asif Ekbal, and Pushpak Bhattacharyya. 2019. Multilingual unsupervised NMT using shared encoder and language-specific decoders. In *ACL*, pages 3083–3089.
- Sennrich, Rico and Biao Zhang. 2019. Revisiting low-resource neural machine translation: A case study. In *ACL*, pages 211–221.
- Sennrich, Rico, Barry Haddow, and Alexandra Birch. 2016a. Improving neural machine translation models with monolingual data. In *ACL*, pages 86–96.
- Sennrich, Rico, Barry Haddow, and Alexandra Birch. 2016b. Neural machine translation of rare words with subword units. In *ACL*, pages 1715–1725.
- Søgaard, Anders, Sebastian Ruder, and Ivan Vulić. 2018. On the limitations of unsupervised bilingual dictionary induction. In *ACL*, pages 778–788.
- Song, Kaitao, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. In *ICML*, pages 5926–5936.
- Srivastava, Nitish, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958.
- Sun, Haipeng, Rui Wang, Kehai Chen, Masao Utiyama, Eiichiro Sumita, and Tiejun Zhao. 2019. Unsupervised bilingual word embedding agreement for unsupervised neural machine translation. In *ACL*, pages 1235–1245.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*, pages 5998–6008.
- Yang, Zhen, Wei Chen, Feng Wang, and Bo Xu. 2018. Unsupervised neural machine translation with weight sharing. In *ACL*, pages 46–55.