

Unsupervised Neural Machine Translation with SMT as Posterior Regularization

Shuo Ren,^{1*} Zhirui Zhang,² Shujie Liu,³ Ming Zhou,³ Shuai Ma¹

¹SKLSDE Lab, Beihang University ¹Beijing Advanced Innovation Center for Big Data and Brain Computing, China

²University of Science and Technology of China, Hefei, China Ming Zhou,³Microsoft Research Asia

¹{shuoren,mashuai}@buaa.edu.cn ²zrustc11@gmail.com ³{shujliu,mingzhou}@microsoft.com

Abstract

Without real bilingual corpus available, unsupervised Neural Machine Translation (NMT) typically requires pseudo parallel data generated with the back-translation method for the model training. However, due to weak supervision, the pseudo data inevitably contain noises and errors that will be accumulated and reinforced in the subsequent training process, leading to bad translation performance. To address this issue, we introduce phrase based Statistic Machine Translation (SMT) models which are robust to noisy data, as posterior regularizations to guide the training of unsupervised NMT models in the iterative back-translation process. Our method starts from SMT models built with pre-trained language models and word-level translation tables inferred from cross-lingual embeddings. Then SMT and NMT models are optimized jointly and boost each other incrementally in a unified EM framework. In this way, (1) the negative effect caused by errors in the iterative back-translation process can be alleviated timely by SMT filtering noises from its phrase tables; meanwhile, (2) NMT can compensate for the deficiency of fluency inherent in SMT. Experiments conducted on *en-fr* and *en-de* translation tasks show that our method outperforms the strong baseline and achieves new state-of-the-art unsupervised machine translation performance.

1 Introduction

Recent years have witnessed the rise and success of Neural Machine Translation (NMT) (Sutskever, Vinyals, and Le 2014; Bahdanau, Cho, and Bengio 2014; Luong, Pham, and Manning 2015; Wu et al. 2016; Vaswani et al. 2017; Hassan et al. 2018). However, NMT relies heavily on large in-domain parallel data, resulting in poor performance on low-resource language pairs (Koehn and Knowles 2017). For some low-resource pairs without any bilingual corpus, how to train NMT models with only a monolingual corpus is a popular and interesting topic.

Existing methods for unsupervised machine translation (Artetxe et al. 2017; Lample, Denoyer, and Ranzato 2017; Yang et al. 2018; Lample et al. 2018) are mainly the modifications of encoder-decoder schema. In their work, source

*The first two authors contributed equally to this work. This work is supported in part by NSFC U1636210 and 61421003, and Shenzhen Institute of Computing Sciences.
Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Pseudo training data (fr-en)	[src1]: ... de sa grand - mère blanche gravement malade , ...
	[trg1]: ... of his severely ill - fated white grandmother ...
	[src2]: ... au chevet de son épouse gravement malade ...
Test Result	[trg2]: ... at his severely ill wife 's bedside ...
	[src3]: ... ils sont malades ou couverts de moisissures !
	[trg3]: ... they are sick or covered with mold !
	...
	[src]: ... avec lui jusqu'à ce qu'il devienne vieux, malade .
	[sys]: ... with him until he's become old, ill-fated .
	[ref]: ... with him until he became old and ill .

Figure 1: The effect of noisy training data. The first training sample contains the noise (“malade” in French means “ill”, not “ill-fated”), leading to the wrong test result (sys).

sentences in two languages are mapped into the same latent space with a shared encoder, which is expected to be the internal information representation irrelevant to the languages themselves. From that target sentences are generated by a shared or different decoders. Some of them also use denoising auto-encoders (Vincent et al. 2010) and adversarial training. Despite the differences in structures and training methods, they reach a consensus to use the pseudo parallel data generated iteratively with the back-translation method (Sennrich, Haddow, and Birch 2016; Zhang et al. 2018a) to train their unsupervised NMT models, i.e. they use monolingual data in the target language and a target-to-source translation model to generate source sentences, then use the pseudo parallel data of generated sources and real targets to train the source-to-target model, and vice versa.

However, since the pseudo data are generated by unsupervised models, random errors and noises are inevitably introduced, such as redundant or unaligned words deviating from the meaning of source sentences. Due to the lack of supervision, those infrequent errors will be accumulated and reinforced by NMT models into frequent patterns during the training iterations, leading to bad translation performance. For instance in Figure 1, the French word “malade” is mistakenly translated into the English word “ill-fated” in the first training sample. With strong abilities to identify and memorize patterns, NMT models mistakenly translate this word into “ill-fated” when “old” (similar to “grandmother”

in the first training sample) occurs in the test. Even so, there are also many good translation patterns (such as “malade” → “ill” or “sick” in the second and third training samples), which could have been extracted in time to guide the NMT models into the correct training direction. The extraction and guidance can be well carried out by Statistical Machine Translation (SMT). As is pointed out by Khayrallah and Koehn (2018), SMT performs better than NMT in tackling noisy data by constructing a strong phrase table with good and frequent translation patterns and filtering out infrequent errors and noises. This gives the motivation that if we incorporate SMT in the training process, unsupervised NMT could benefit from the robustness of SMT to noisy data.

In this paper, we propose to leverage SMT to denoise and guide the training of unsupervised NMT models in the iterative back-translation process. Different from previous work (He et al. 2016; Tang et al. 2016; Wang et al. 2017) introducing SMT into NMT by changing model structures in supervised scenarios, we adopt the framework of posterior regularization (Ganchev et al. 2010) to leave model structures unchanged. Our method starts from initial SMT models built with pre-trained language models and word-level translation tables inferred from cross-lingual embeddings. Then SMT models and NMT models are trained jointly in a unified Expectation Maximization (EM) training framework. In each iteration, as desired distributions, SMT models are expected to correct NMT models timely with denoised pseudo data generated in a constrained search space of reliable translation patterns. Based on that, enhanced NMT models can generate better pseudo data for SMT to extract phrases of higher quality, so that they can benefit from each other incrementally. In this way, infrequent errors in NMT models can be eliminated with the constraints exerted by SMT features, while NMT can compensate for the deficiency in smoothness inherent in SMT models. Experiments conducted on en-fr and en-de translation tasks show that our method significantly outperforms the strong baseline (Lample et al. 2018) and achieves the new state-of-the-art translation performance in unsupervised machine translation.

2 Background

2.1 Neural Machine Translation

Given a source sentence $\mathbf{x} = (x_1, x_2, \dots, x_l)$ and a target one $\mathbf{y} = (y_1, y_2, \dots, y_m)$, Neural Machine Translation (NMT) directly models the word-level translation probability with parameters θ as:

$$p(y_i | \mathbf{x}, \mathbf{y}_{<i}; \theta) = \text{softmax}(g(\mathbf{h}_{y_i}, \mathbf{h}_{\mathbf{y}_{<i}}, \mathbf{c}_i; \theta)) \quad (1)$$

in which $g(\cdot)$ denotes a non-linear function extracting features to predict the target word y_i from the decoder states (\mathbf{h}_{y_i} and $\mathbf{h}_{\mathbf{y}_{<i}}$) and the context vector \mathbf{c}_i calculated with the encoder and attention mechanism. Then the sentence-level translation probability $p(\mathbf{y} | \mathbf{x}; \theta)$ is calculated by $p(\mathbf{y} | \mathbf{x}; \theta) = \prod_{i=1}^m p(y_i | \mathbf{x}, \mathbf{y}_{<i}; \theta)$. As for training, given a parallel corpus $\{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, the objective function is to maximize $\log p(\mathbf{y}_n | \mathbf{x}_n; \theta)$ over the whole training set.

2.2 Phrase-based Statistic Machine Translation

The current approach of Statistic Machine Translation (SMT) is typically based on the log-linear model proposed by Och and Ney (2002). According to it, the translation probability from sentence \mathbf{x} to sentence \mathbf{y} is formulated as:

$$p(\mathbf{y} | \mathbf{x}; \lambda_1^M) = \frac{\exp[\sum_{m=1}^M \lambda_m h_m(\mathbf{x}, \mathbf{y})]}{\sum_{\tilde{\mathbf{y}}} \exp[\sum_{m=1}^M \lambda_m h_m(\mathbf{x}, \tilde{\mathbf{y}})]} \quad (2)$$

where $h_m(\mathbf{x}, \mathbf{y}) = \log \phi_m(\mathbf{x}, \mathbf{y})$ denotes the m^{th} feature.

In phrase based SMT (PBSMT) (Koehn, Och, and Marcu 2003), the sentence pair is segmented into a sequence of phrases $\tilde{\mathbf{x}}_1^I$ and $\tilde{\mathbf{y}}_1^J$, where I and J are the counts of phrases. During training, given a bilingual corpus, PBSMT first infers word alignment, based on which phrase pairs are derived and stored in the phrase table, as well as translation probabilities. Other features such as a distortion model can also be learned with the extracted phrase pairs. The feature weights λ_1^M can be optimized by MERT (Och 2003) with a validation set. During decoding, PBSMT generates translation candidates $\tilde{\mathbf{y}}$ bottom up via the CKY algorithm, ranked with scores given by the log-linear model in Eq.(2).

2.3 Posterior Regularization

Posterior regularization (Ganchev et al. 2010) is a framework for structured, weakly supervised learning, which incorporates indirect supervision from a desired distribution $q(\mathbf{y})$ via constraints on posterior distribution $p(\mathbf{y} | \mathbf{x}_n; \theta)$ imposed by a Kullback-Leible (KL) divergence as follows:

$$F(q; \theta) = \mathcal{L}(\theta) - \sum_{n=1}^N \min_{q \in Q} \text{KL}(q(\mathbf{y}) || p(\mathbf{y} | \mathbf{x}_n; \theta)) \quad (3)$$

where $\mathcal{L}(\theta)$ is the original likelihood of model $p(\mathbf{y} | \mathbf{x}; \theta)$, and Q is a constraint posterior set satisfying:

$$Q = \{q(\mathbf{y}) : \mathbf{E}_q[\phi(\mathbf{x}, \mathbf{y})] \leq \mathbf{b}\} \quad (4)$$

in which constraints features $\phi(\mathbf{x}, \mathbf{y})$ are bounded by \mathbf{b} .

To maximize $F(q; \theta)$, Ganchev et al. (2010) propose an EM framework (McLachlan and Krishnan 2007) as:

$$\begin{aligned} E : q^{t+1} &= \arg \min_{q \in Q} \text{KL}(q(\mathbf{y}) || p(\mathbf{y} | \mathbf{x}_n; \theta^t)) \\ M : \theta^{t+1} &= \arg \max_{\theta} \mathcal{L}(\theta) + \mathbf{E}_{q^{t+1}}[\log p(\mathbf{y} | \mathbf{x}_n; \theta)] \end{aligned} \quad (5)$$

However, there may be a problem as pointed out by Zhang et al. (2017) that it is hard to set a reasonable bound \mathbf{b} if we directly apply posterior regularization to NMT. To solve this problem, we follow their practice of representing the desired distribution $q(\mathbf{y})$ as the log-linear model described in Eq.(2). In this way, SMT models directly act as the posterior regularization to constrain NMT models $p(\mathbf{y} | \mathbf{x}_n; \theta^t)$.

3 Method

3.1 Overview

Due to the lack of supervision, noises and infrequent errors in the pseudo data generated by unsupervised NMT models will be accumulated and reinforced in the iterative back-translation process (shown in the shadow area in Figure 2).

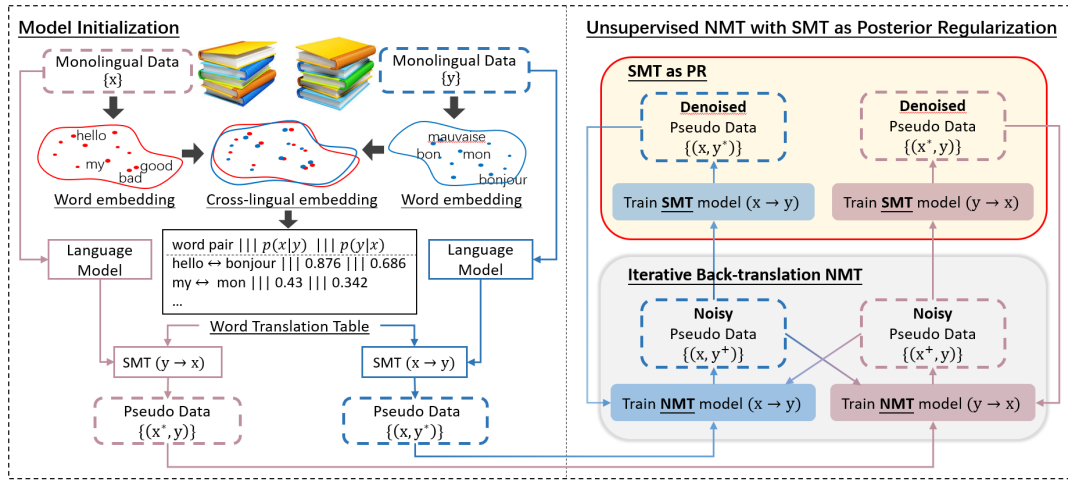


Figure 2: Method overview. The whole procedure mainly consists of two parts as the left and the right.

To address this issue, we introduce SMT as posterior regularization (the red frame above that) to denoise and guide the training of NMT, thus the noises being eliminated timely.

The whole procedure of our method mainly consists of two parts shown in the left and right of Figure 2. Given a language pair X-Y, for **model initialization**, we build two initial SMT models with language models pre-trained using monolingual data, and word translation tables inferred from cross-lingual embeddings according to the approach in 3.2. Then the initial SMT models will generate pseudo data to warm up two NMT models. Note that the NMT models are trained using not only the pseudo data generated by SMT models, but those generated by reverse NMT models with the **iterative back-translation** method. After that, the NMT-generated pseudo data are fed to SMT models. As **posterior regularization (PR)**, SMT models timely filter out noises and infrequent errors by constructing strong phrase tables with good and frequent translation patterns, and then generate denoised pseudo data to guide the subsequent NMT training. Benefiting from that, NMT then produces better pseudo data for SMT to extract phrases of higher quality, meanwhile compensating for the deficiency in smoothness inherent in SMT via back-translation. Those two steps are unified in the EM framework described in 3.3, where NMT and SMT models are trained jointly and boost each other incrementally until final convergence.

3.2 Initialization

Our initial SMT models are built with word-based phrase tables and two pre-trained language models via Moses¹. For the word translation table, we first train word embeddings using monolingual corpora for two languages respectively. Based on that, we adopt the method proposed by Artetxe et al. (2018) to obtain respective cross-lingual embeddings $\{e_{x_i}\}_{i=1}^S$ and $\{e_{y_j}\}_{j=1}^T$, where S and T are vocabulary sizes.

¹<https://github.com/moses-smt/mosesdecoder>

Then the word translation probability from word x_i to y_j is:

$$p(y_j|x_i) = \frac{\exp[\lambda \cos(e_{x_i}, e_{y_j})]}{\sum_k \exp[\lambda \cos(e_{x_i}, e_{y_k})]} \quad (6)$$

where λ is a hyper-parameter to control the peakiness of the distribution. The calculation of $p(x_i|y_j)$ is similar to Eq.(6). Based on the above, we choose top- k translation candidates for each word in our initial phrase table. We only use two features in our initial phrase tables, i.e. translation probabilities and inverse translation probabilities.

3.3 Unsupervised NMT with SMT as PR

As is mentioned in 3.1, SMT plays a role in denoising and is leveraged as posterior regularization for NMT. Therefore, we replace the posterior regularization term $q(\mathbf{y})$ in Eq.(3) with the SMT models ($\mathbf{x} \rightarrow \mathbf{y}$) and ($\mathbf{y} \rightarrow \mathbf{x}$) in Figure 2, which will be denoted by $\vec{p}_s(\mathbf{y}|\mathbf{x})$ and $\overleftarrow{p}_s(\mathbf{x}|\mathbf{y})$. By the way, the NMT models ($\mathbf{x} \rightarrow \mathbf{y}$) and ($\mathbf{y} \rightarrow \mathbf{x}$) will be denoted by $\vec{p}_n(\mathbf{y}|\mathbf{x}; \theta_{\mathbf{x} \rightarrow \mathbf{y}})$ and $\overleftarrow{p}_n(\mathbf{x}|\mathbf{y}; \theta_{\mathbf{x} \leftarrow \mathbf{y}})$, where $\theta_{\mathbf{x} \rightarrow \mathbf{y}}$ and $\theta_{\mathbf{x} \leftarrow \mathbf{y}}$ are parameters. Then, given monolingual corpora $\{\mathbf{x}_i\}_{i=1}^M$ and $\{\mathbf{y}_j\}_{j=1}^N$, we formulate the training objective as:

$$\begin{aligned} \mathcal{J}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}, \vec{p}_s, \overleftarrow{p}_s) &= \bar{\mathcal{L}}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}) \\ &- \sum_{i=1}^M \min_{\vec{p}_s} \mathbf{KL}(\vec{p}_s(\mathbf{y}|\mathbf{x}_i) || \vec{p}_n(\mathbf{y}|\mathbf{x}_i; \theta_{\mathbf{x} \rightarrow \mathbf{y}})) \\ &- \sum_{j=1}^N \min_{\overleftarrow{p}_s} \mathbf{KL}(\overleftarrow{p}_s(\mathbf{x}|\mathbf{y}_j) || \overleftarrow{p}_n(\mathbf{x}|\mathbf{y}_j; \theta_{\mathbf{x} \leftarrow \mathbf{y}})) \end{aligned} \quad (7)$$

where $\bar{\mathcal{L}}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}})$ corresponds to the training objective of iterative back-translation for NMT models, which is

$$\begin{aligned} &\bar{\mathcal{L}}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}) \\ &= \sum_{i=1}^M \mathbf{E}_{\mathbf{y} \sim \vec{p}_n(\mathbf{y}|\mathbf{x}_i; \theta_{\mathbf{x} \rightarrow \mathbf{y}})} [\log \overleftarrow{p}_n(\mathbf{x}_i|\mathbf{y}; \theta_{\mathbf{x} \leftarrow \mathbf{y}})] \\ &+ \sum_{j=1}^N \mathbf{E}_{\mathbf{x} \sim \overleftarrow{p}_n(\mathbf{x}|\mathbf{y}_j; \theta_{\mathbf{x} \leftarrow \mathbf{y}})} [\log \vec{p}_n(\mathbf{y}_j|\mathbf{x}; \theta_{\mathbf{x} \rightarrow \mathbf{y}})] \end{aligned} \quad (8)$$

and two Kullback-Leibler divergence (KL) terms denote the posterior regularizations for two NMT models respectively.

Based on that, the training processes of iterative back-translation for NMT and SMT models as posterior regularization are unified into a single objective \mathcal{J} . Then, we modulate the EM algorithm in Eq.(5) to optimize it as follows:

$$\begin{aligned}
E : \overleftarrow{p}_s^{t+1} &= \arg \max_{\overleftarrow{p}_s} \mathcal{J}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}, \overrightarrow{p}_s, \overleftarrow{p}_s) \\
&= \arg \min_{\overleftarrow{p}_s} \mathbf{KL}(\overleftarrow{p}_s(\mathbf{x}|\mathbf{y}_j) || \overleftarrow{p}_n(\mathbf{x}|\mathbf{y}_j; \theta_{\mathbf{x} \leftarrow \mathbf{y}}^t)) \\
\overrightarrow{p}_s^{t+1} &= \arg \max_{\overrightarrow{p}_s} \mathcal{J}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}, \overrightarrow{p}_s, \overleftarrow{p}_s) \\
&= \arg \min_{\overrightarrow{p}_s} \mathbf{KL}(\overrightarrow{p}_s(\mathbf{y}|\mathbf{x}_i) || \overrightarrow{p}_n(\mathbf{y}|\mathbf{x}_i; \theta_{\mathbf{x} \rightarrow \mathbf{y}}^t)) \\
M : \theta_{\mathbf{x} \leftarrow \mathbf{y}}^{t+1} &= \arg \max_{\theta_{\mathbf{x} \leftarrow \mathbf{y}}} \mathcal{J}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}, \overrightarrow{p}_s, \overleftarrow{p}_s) \\
&= \arg \max_{\theta_{\mathbf{x} \leftarrow \mathbf{y}}} \{ \mathbf{E}_{\overleftarrow{p}_s^{t+1}} [\log \overleftarrow{p}_n(\mathbf{x}|\mathbf{y}_j; \theta_{\mathbf{x} \leftarrow \mathbf{y}})] \\
&\quad + \mathbf{E}_{\overrightarrow{p}_n(\mathbf{y}|\mathbf{x}_i; \theta_{\mathbf{x} \rightarrow \mathbf{y}}^t)} [\log \overleftarrow{p}_n(\mathbf{x}_i|\mathbf{y}; \theta_{\mathbf{x} \leftarrow \mathbf{y}})] \} \\
\theta_{\mathbf{x} \rightarrow \mathbf{y}}^{t+1} &= \arg \max_{\theta_{\mathbf{x} \rightarrow \mathbf{y}}} \mathcal{J}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}, \overrightarrow{p}_s, \overleftarrow{p}_s) \\
&= \arg \max_{\theta_{\mathbf{x} \rightarrow \mathbf{y}}} \mathbf{E}_{\overrightarrow{p}_s^{t+1}} [\log \overrightarrow{p}_n(\mathbf{y}|\mathbf{x}_i; \theta_{\mathbf{x} \rightarrow \mathbf{y}})] \\
&\quad + \mathbf{E}_{\overleftarrow{p}_n(\mathbf{x}|\mathbf{y}_j; \theta_{\mathbf{x} \leftarrow \mathbf{y}}^t)} [\log \overrightarrow{p}_n(\mathbf{y}_j|\mathbf{x}; \theta_{\mathbf{x} \rightarrow \mathbf{y}})]
\end{aligned} \tag{9}$$

Briefly speaking, in the E-step, we optimize the desired distributions represented by SMT to minimize the KL distance between SMT models and NMT models. In the M-step, we optimize NMT models using the pseudo data generated by SMT models and the corresponding reverse NMT models to fit the desired distributions and meanwhile perform back-translation iterations. We will give the specific equation for updating parameters in 3.4.

3.4 Training Algorithm

We combine the model initialization and the whole training procedure into Algorithm 1 as follows.

According to Eq.(9), in the E-step, we need to minimize the gap between SMT models and NMT models. However, this step cannot be done by traditional gradient descent methods. Approximately, we train SMT models using the pseudo data generated by the corresponding NMT models to fit the mode of NMT posterior distributions. Thus the KL divergence between them is diminished. This step corresponds to the 7th and 8th lines in Algorithm 1, meaning SMT extracts good and frequent translation patterns from the data generated by current NMT models to finish denoising.

In the M-step, we optimize two NMT models with gradient descent methods. We formulate the updating for $\theta_{\mathbf{x} \leftarrow \mathbf{y}}$ in Eq.(10), to which that for $\theta_{\mathbf{x} \rightarrow \mathbf{y}}$ is similar.

$$\begin{aligned}
&\nabla_{\theta_{\mathbf{x} \leftarrow \mathbf{y}}} \mathcal{J}(\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}, \overrightarrow{p}_s, \overleftarrow{p}_s) \\
&= \mathbf{E}_{\mathbf{x} \sim \overleftarrow{p}_s(\mathbf{x}|\mathbf{y}_j)} \nabla_{\theta_{\mathbf{x} \leftarrow \mathbf{y}}} \log \overleftarrow{p}_n(\mathbf{x}|\mathbf{y}_j; \theta_{\mathbf{x} \leftarrow \mathbf{y}}) \\
&\quad + \mathbf{E}_{\mathbf{y} \sim \overrightarrow{p}_n(\mathbf{y}|\mathbf{x}_i; \theta_{\mathbf{x} \rightarrow \mathbf{y}})} \nabla_{\theta_{\mathbf{x} \leftarrow \mathbf{y}}} \log \overleftarrow{p}_n(\mathbf{x}_i|\mathbf{y}; \theta_{\mathbf{x} \leftarrow \mathbf{y}})
\end{aligned} \tag{10}$$

Algorithm 1: Unsupervised NMT with SMT as PR

Input: Monolingual data $X = \{\mathbf{x}_i\}_{i=1}^M$ and $Y = \{\mathbf{y}_j\}_{j=1}^N$
Output: Parameters of two NMT models: $\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}$

- 1 Train language models l_x and l_y using X and Y
- 2 Infer word translation tables t_{xy} and t_{yx} as in 3.2
- 3 $t := 0$
- while not convergence do**
- 4 Sample data $\{\mathbf{x}_t\} \in X$ and $\{\mathbf{y}_t\} \in Y$
- 5 // **E-step:**
- 6 **if** $t = 0$ **then**
- 7 Initialize \overrightarrow{p}_s^0 and \overleftarrow{p}_s^0 using l_x, l_y, t_{xy} and t_{yx}
- 8 **else**
- 9 Generate pseudo data $\{(\mathbf{x}_t, \mathbf{y}_t^+)\}$ and $\{(\mathbf{x}_t^+, \mathbf{y}_t)\}$
- 10 using models \overrightarrow{p}_n^t and \overleftarrow{p}_n^t respectively
- 11 Train \overrightarrow{p}_s^t and \overleftarrow{p}_s^t using $(\mathbf{x}_t, \mathbf{y}_t^+)$ and $(\mathbf{x}_t^+, \mathbf{y}_t)$
- 12 Generate pseudo data $\{(\mathbf{x}_t, \mathbf{y}_t^+)\}$ and $\{(\mathbf{x}_t^+, \mathbf{y}_t)\}$ using
- 13 \overrightarrow{p}_n^t and \overleftarrow{p}_n^t respectively
- 14 Train \overrightarrow{p}_n^t and \overleftarrow{p}_n^t using $\{(\mathbf{x}_t^+, \mathbf{y}_t)\} \cup \{(\mathbf{x}_t, \mathbf{y}_t^+)\}$ and
- 15 $\{(\mathbf{x}_t, \mathbf{y}_t^+)\} \cup \{(\mathbf{x}_t^+, \mathbf{y}_t)\}$
- 16 $t := t + 1$
- 17 **return** $\theta_{\mathbf{x} \rightarrow \mathbf{y}}, \theta_{\mathbf{x} \leftarrow \mathbf{y}}$

This step corresponds to lines 14 to 17 in Algorithm 1. A difficulty here is the exponential search space of the translation candidates. To address it, we leverage the sampling method (Shen et al. 2015) and simply generate the top target sentence for approximation in our experiments. Note that in the 11th line, NMT models are trained using the denoised pseudo data generated by SMT models only, while in the 13th line, the mixed data of those and the pseudo data generated by the reverse NMT models are used. The intention here is to first use the denoised pseudo data to correct the NMT models established before, and then apply iterative back-translation to boost NMT models under the guide of the denoised data. NMT also makes up for the deficiency in smoothness of SMT in this step. In this way, SMT and NMT models can benefit from each other in the EM iterations.

4 Experiments

4.1 Setup

Dataset In our experiments, we consider two language pairs, English-French and English-German. For each language, we use 50 million monolingual sentences in NewsCrawl, a monolingual dataset from WMT, which is the same as the previous work (Artetxe et al. 2017; Lample et al. 2018). For the convenience of comparison, we use *newstest* 2014 as the test set for the English-French pair, and *newstest* 2014 as well as *newstest* 2016 for the English-German pair.

Preprocess We use Moses scripts for word tokenization and truecasing. In model initialization, we use the public im-

Method	fr-en	en-fr	de-en (2014)	en-de (2014)	de-en (2016)	en-de (2016)
(Artetxe et al. 2017)	15.56	15.13	10.21	6.89	-	-
(Lample, Denoyer, and Ranzato 2017)	14.31	15.05	-	-	13.33	9.64
(Yang et al. 2018)	15.58	16.97	-	-	14.62	10.86
(Lample et al. 2018), NMT	24.18	25.41	-	-	21.00	17.16
(Lample et al. 2018), PBSMT	27.16	28.11	-	-	22.68	17.77
(Lample et al. 2018), NMT+PBSMT	26.29	27.12	-	-	22.06	17.52
(Lample et al. 2018), PBSMT+NMT	27.68	27.60	-	-	25.19	20.23
Our Method	28.79	29.21	20.04	16.43	25.92	21.07
(+ R2L regularization)	28.92	29.53	20.43	16.97	26.32	21.65

Table 1: Comparison with previous methods.

plementation of word2vec² to train monolingual word embeddings of each language, and vecmap³ to obtain cross-lingual embeddings of both language pairs. For NMT, we use the modified version of the public implementation⁴ of Transformer (Vaswani et al. 2017). We share the vocabulary space of 50,000 BPE codes (Sennrich, Haddow, and Birch 2015) for source and target languages. For each language pair, we train two independent NMT models for different translation directions (i.e., source to target and target to source) with shared embedding layers of source and target sides. For SMT, we use the Moses implementation of PBSMT systems with Salm (Johnson et al. 2007), which can denoise and reduce the size of phrase tables. And we use the default features defined in Moses for our PBSMT models.

4.2 Comparison

Baselines Our proposed method is compared with four baselines of unsupervised machine translation listed in the upper area of Table 1, among which the fourth baseline contains several methods. Given a language pair, the first two baselines (Artetxe et al. 2017; Lample, Denoyer, and Ranzato 2017) use a shared encoder and different decoders for the two languages. The third baseline (Yang et al. 2018) uses different encoders and decoders, and introduces a weight sharing mechanism. The fourth baseline (Lample et al. 2018) uses a shared encoder and decoder in their NMT systems. As for the training method, the second and third baselines use adversarial training. All of the four baselines use denoising auto-encoder and iterative back-translation.

Note that the fourth baseline contains four methods. “NMT” means unsupervised NMT models, while “PBSMT” denotes unsupervised SMT models with the back-translation method performed by SMT. “NMT+PBSMT” and “PBSMT+NMT” simply combine the best pseudo data that the former generates into the final iteration of the latter. Different from our proposed method, the training processes of NMT and SMT models in their methods are independent.

Results and Discussion The comparison results are reported in Table 1. The BLEU scores are calculated by *multi-bleu.pl*. From the table, we find that our method significantly

outperforms all the baselines even the strong one (Lample et al. 2018). We elaborate the reasons as follows.

(1) Our proposed method significantly improves the performance over the “NMT” and “PBSMT” of (Lample et al. 2018). This is because unsupervised NMT methods suffer from the noise problem while PBSMT is inherently deficient in fluency just as the case study in 4.5 shows. Our method can compensate for the deficiencies of them by combining the training processes of them. (2) Notice that “NMT+PBSMT” performs even worse than pure “PBSMT”, which may be caused by accumulated errors in the iterations of NMT models. Due to the lack of timely denoising methods, infrequent errors and noises are repeated and reinforced as frequent ones by unsupervised NMT, so that even PBSMT could not distinguish them from good patterns in the last iteration. (3) The performance gained by “PBSMT+NMT” verifies combining data of high quality into NMT training could be a better choice. But the simple combination in their method is not able to make the best of both models. In their method, NMT and SMT models are trained independently so that the bad patterns within the models themselves cannot be well removed due to weak supervision. In contrast, our proposed method integrates the training of NMT and SMT models in a unified EM framework where they can boost each other incrementally. The noises and errors generated by NMT models can be reduced in time by SMT as posterior regularization, while NMT can compensate for the deficiency of smoothness inherent in SMT models. Therefore, our proposed method still outperforms “PBSMT+NMT”.

Apart from SMT as posterior regularization, our framework can be easily extended to incorporate other posterior regularization methods without changing model structures, such as the target-bidirectional agreement regularization (Zhang et al. 2018b). This regularization can help deal with the problem of exposure bias in supervised NMT, where another “reversed” NMT model is trained using data of reversed sentences from left to right. Then the “reversed” NMT model is leveraged to generate pseudo data for training the original NMT model. Specifically, we introduce the R2L regularization after the final training iteration of NMT models (i.e., NMT2 in Table 2). With this extension, we achieve higher performance (+R2L regularization in Table 1).

²<https://github.com/tmikolov/word2vec>

³<https://github.com/artetxem/vecmap>

⁴<https://github.com/tensorflow/tensor2tensor>

Steps	fr-en	en-fr	de-en	en-de	ave
E-step (SMT0)	15.34	11.74	11.03	8.14	11.56
M-step (NMT0)	24.06	24.82	16.29	12.88	+7.95
E-step (SMT1)	26.49	27.64	17.34	14.81	+2.06
M-step (NMT1)	28.29	29.02	19.61	16.02	+1.67
E-step (SMT2)	28.64	29.21	19.87	16.29	+0.23
M-step (NMT2)	28.79	29.17	20.04	16.43	+0.11

Table 2: Test BLEU on *newstest* 2014 in different steps.

4.3 Model Evolution

We conduct several EM iterations in our experiments, and record the test BLEU scores on *newstest* 2014 after each E-step (SMT) and M-step (NMT) in Table 2. We have tried more steps but the models do not converge after three EM iterations. For the convenience of comparison, in the last column of the table, we also list the average improvement of four translation models after each step. From the table, first, we find NMT and SMT models improve incrementally after each iteration, which accords with our proposed motivation. Note that the improvements between adjacent NMT steps are exactly contributions made by SMT as posterior regularization. Second, the models improve the most in the first EM iteration and nearly converge at the third EM iteration.

Additionally, we also compare the translation performance on sentences of different lengths as iteration steps progress. We group the sentences in the fr-en test set by length as shown by the three curves in Figure 3. Then, we record the BLEU scores of different groups after each step. From the figure, we find the models converge much slower on longer sentences, which indicates that it is easier for the models to learn shorter sentences.




Figure 3: Test BLEU on sentences grouped by length.

4.4 Discussion on Initialization

In this subsection, we delve into the initialization stage which is crucial to our method. In that stage, there are three hyper parameters described in 3.2 that should be taken into account, i.e., the peakiness controller λ , the vocabulary size S or T , and the number of translation candidates k for each word. Since the performance of initialization can be evaluated by SMT0, we adjust the hyper-parameters and measure the fr-en test BLEU of SMT0 models accordingly. For brevity, we let $S = T = V$ in our experiments. The results are illustrated in Figure 4. From this figure, we find




Figure 4: Test of initial models with various hyper-params.

that k and V have much bigger impacts on the initial model SMT0 than λ . With the value of λ increasing, the performance of SMT0 gradually improves but starts to decline a bit after around 20. This is because the larger λ will make the distribution in Eq.(6) sharper, severely restricting the search spaces of SMT models. Similarly, the performance of SMT0 improves in accord with the value of k or V going up. But the improvement stops after certain thresholds (about 80 of k and 50000 of V). The reason may be the useful information provided by word-translation tables is saturated after those.

We also tried other initialization methods in our experiments, such as directly using the pseudo parallel data constructed from word-by-word translation to warm up NMT models. We compare NMT0 models warmed up with this method (without SMT0) to NMT0 in our proposed method (with SMT0) in the following table, which stresses the necessity of SMT0 and the importance of good initialization.

Initialization Method	fr-en	en-fr	de-en	en-de
NMT0 without SMT0	12.29	12.46	7.32	4.81
NMT0 with SMT0	24.06	24.82	16.29	12.88

Table 3: The necessity of SMT0 in model initialization. The numbers in this table are BLEU scores on *newstest* 2014.

4.5 Case Study

To further demonstrate the effectiveness of our method, we select some cases from translation results (fr-en) and compare the translations generated by models of different training steps. The results are listed in Table 4. In the first case, which is exactly the example in the Introduction, the word “malade” in French is wrongly translated into “ill-fated” in English by NMT0. As we can see, this error has been corrected in NMT1 after the guidance of SMT1. In the second case, apart from the wrongly aligned word “bâtisse-là” to “canopy-back business” by NMT1, there is also a redundant phrase “plenty of” generated by it. Those errors are both corrected after the regularization of SMT1. In the third case, we also reach the same conclusion that NMT1 can benefit from SMT1 and rectify the mistake on “rendu visite à”. There is also an interesting phenomenon from case three of NMT adhering to “from” which makes the sentence more fluent, even though this word is missed by SMT models. In a word, the above analysis verifies that noises and errors in unsupervised NMT models can be eliminated timely by SMT models as posterior regularization with our method.

Source	J'ai eu des relations difficiles avec lui jusqu'à ce qu'il devienne vieux, <u>malade</u> .
SMT0	I've gotten of difficult relations with him until he will become old, <u>sick</u> .
NMT0	I've had difficult relations with him until he's become old, <u>ill-fated</u> .
SMT1	I've had difficult relationships with him until he became old, <u>sick</u> .
NMT1	I had difficult relations with him until he became old and <u>sick</u> .
Reference	I had a difficult relationship with him until he became old and <u>ill</u> .
Source	Le fonds d'investissement qui était propriétaire de cette <u>bâtisse-là</u> avait des choix à faire.
SMT0	The owner of this <u>underline</u> building, so had to make a choice of which was an investment fund.
NMT0	The investment fund that was an owner of that <u>canopy-back business</u> had <u>plenty</u> of choice to do.
SMT1	The investment fund that was the owner of this <u>building</u> just had to make choices.
NMT1	The investment fund that was the owner of this <u>building</u> had choices to make.
Reference	The investment fund that owned the <u>building</u> had to make a choice.
Source	M. Dutton a <u>rendu visite</u> à Mme Plibersek pour garantir qu'aucun dollar <u>du</u> plan de sauvetage ne sera dépensé en bureaucratie supplémentaire.
SMT0	Mr Dutton <u>paid a visit to</u> Ms Plibersek to guarantee that the greenback no rescue plan of not be spent in extra bureaucracy.
NMT0	Mr Dutton <u>said</u> Ms Plibersek' visit to guarantee any dollar <u>from</u> the rescue plan will be spent in extra bureaucracy.
SMT1	Mr Dutton was <u>visiting</u> Ms Plibersek to guarantee that no dollar rescue plan will be spent on additional bureaucracy.
NMT1	Mr Dutton <u>paid a visit to</u> Ms Plibersek to guarantee that no dollar <u>from</u> the rescue plan will be spent on extra bureaucracy.
Reference	Mr Dutton <u>called on</u> Ms Plibersek to guarantee that not one dollar <u>out of</u> the rescue package would be spent on additional bureaucracy.

Table 4: Cases of translation results from French to English in *newstest* 2014. The models of SMT0, NMT0, SMT1 and NMT1 are corresponding to the steps in Table 2.

From these cases, we find that SMT can also benefit from NMT models. Even though the meanings of the key words could be captured by SMT, the outputs of SMT0 are not fluent especially in the second case. This problem is relieved in SMT1, after SMT is fed with more fluent pseudo data generated by NMT0, which validates that SMT and NMT can incrementally boost each other with our method.

5 Related Work

Previous unsupervised neural machine translation systems (Artetxe et al. 2017; Lample, Denoyer, and Ranzato 2017; Yang et al. 2018) are mainly the modifications of the current encoder-decoder structure. To constrain outputs of encoders for two languages into a same latent space, Artetxe et al. (2017), and Lample et al. (2017) use a shared encoder, while Yang et al. (2018) use a weight sharing mechanism. Denoising auto-encoder (Vincent et al. 2010) and adversarial training methods are also leveraged to improve the ability of encoders. Besides, iterative back-translation is applied to generated pseudo parallel data for cross-lingual training.

After that, Lample et al. (2018) summarize three principles for unsupervised machine translation, which are initialization, language modeling and iterative back-translation, and propose some effective methods with simplified training procedures. Four methods are leveraged in their work, including unsupervised NMT, unsupervised PBSMT and two combinations of them. Our method is different from them. In their methods, SMT and NMT are treated as independent models so that they suffer from respective deficiencies and cannot benefit from each other in their training processes. In contrast, we combine them into a unified EM training framework and enable them to improve jointly and boost each other incrementally, where NMT models are responsible for smoothing and fluency, while SMT models are responsible for denoising and guiding NMT models.

Moreover, there has been some work exploiting SMT features to improve supervised NMT. In He et al. (2016), the probability calculated by NMT is integrated as a feature into a log-linear model. After that, Tang et al. (2016) and Wang et al. (2017) leverage gate mechanisms to introduce a phrase table or candidates provided by SMT into NMT models. Different from them, we leave the model structures unchanged via the framework of posterior regularization. Zhang et al. (2017) also integrate more prior knowledge into the training of NMT with the help of posterior regularization. But there is a major difference that we introduce the successful practice of iterative back-translation into this framework with a unified EM training algorithm, where SMT and NMT models can benefit from each other. Additionally, in unsupervised scenarios, our SMT features are learned from scratch and improved incrementally, rather than pre-trained from real bilingual data and fixed during the whole procedure.

6 Conclusion

In this paper, we introduce SMT models as posterior regularization to denoise and guide unsupervised NMT models with the ability of constructing more reliable phrase tables and eliminating the infrequent and bad patterns generated in the back-translation iterations of NMT. We unify SMT and NMT models within the EM training algorithm where they can be trained jointly and benefit from each other incrementally. In the experiments conducted on en-fr and en-de language pairs, our method significantly outperforms previous methods, and achieves the new state-of-the-art performance of unsupervised machine translation, which demonstrates the effectiveness of our method. In the future, we may delve into the initialization stage, which is crucial to the final performance of the proposed method.

References

- Artetxe, M.; Labaka, G.; Agirre, E.; and Cho, K. 2017. Unsupervised neural machine translation. *arXiv preprint arXiv:1710.11041*.
- Artetxe, M.; Labaka, G.; and Agirre, E. 2018. Generalizing and improving bilingual word embedding mappings with a multi-step framework of linear transformations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 5012–5019.
- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Ganchev, K.; Gillenwater, J.; Taskar, B.; et al. 2010. Posterior regularization for structured latent variable models. *Journal of Machine Learning Research* 11(Jul):2001–2049.
- Hassan, H.; Aue, A.; Chen, C.; Chowdhary, V.; Clark, J.; Federmann, C.; Huang, X.; Junczys-Dowmunt, M.; Lewis, W.; Li, M.; et al. 2018. Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.
- He, W.; He, Z.; Wu, H.; and Wang, H. 2016. Improved neural machine translation with smt features. In *AAAI*, 151–157.
- Johnson, H.; Martin, J.; Foster, G.; and Kuhn, R. 2007. Improving translation quality by discarding most of the phrasetable. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*.
- Khayrallah, H., and Koehn, P. 2018. On the impact of various types of noise on neural machine translation. *arXiv preprint arXiv:1805.12282*.
- Koehn, P., and Knowles, R. 2017. Six challenges for neural machine translation. *arXiv preprint arXiv:1706.03872*.
- Koehn, P.; Och, F. J.; and Marcu, D. 2003. Statistical phrase-based translation. In *Proceedings of the 2003 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology-Volume 1*, 48–54. Association for Computational Linguistics.
- Lample, G.; Ott, M.; Conneau, A.; Denoyer, L.; and Ranzato, M. 2018. Phrase-based & neural unsupervised machine translation. *arXiv preprint arXiv:1804.07755*.
- Lample, G.; Denoyer, L.; and Ranzato, M. 2017. Unsupervised machine translation using monolingual corpora only. *arXiv preprint arXiv:1711.00043*.
- Luong, M.-T.; Pham, H.; and Manning, C. D. 2015. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*.
- McLachlan, G., and Krishnan, T. 2007. *The EM algorithm and extensions*, volume 382. John Wiley & Sons.
- Och, F. J., and Ney, H. 2002. Discriminative training and maximum entropy models for statistical machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*, 295–302. Association for Computational Linguistics.
- Och, F. J. 2003. Minimum error rate training in statistical machine translation. In *Proceedings of the 41st Annual Meeting on Association for Computational Linguistics-Volume 1*, 160–167. Association for Computational Linguistics.
- Sennrich, R.; Haddow, B.; and Birch, A. 2015. Neural machine translation of rare words with subword units. *arXiv preprint arXiv:1508.07909*.
- Sennrich, R.; Haddow, B.; and Birch, A. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 86–96.
- Shen, S.; Cheng, Y.; He, Z.; He, W.; Wu, H.; Sun, M.; and Liu, Y. 2015. Minimum risk training for neural machine translation. *arXiv preprint arXiv:1512.02433*.
- Sutskever, I.; Vinyals, O.; and Le, Q. V. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, 3104–3112.
- Tang, Y.; Meng, F.; Lu, Z.; Li, H.; and Yu, P. L. 2016. Neural machine translation with external phrase memory. *arXiv preprint arXiv:1606.01792*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 6000–6010.
- Vincent, P.; Larochelle, H.; Lajoie, I.; Bengio, Y.; and Manzagol, P.-A. 2010. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. *Journal of Machine Learning Research* 11(Dec):3371–3408.
- Wang, X.; Lu, Z.; Tu, Z.; Li, H.; Xiong, D.; and Zhang, M. 2017. Neural machine translation advised by statistical machine translation. In *AAAI*, 3330–3336.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Yang, Z.; Chen, W.; Wang, F.; and Xu, B. 2018. Unsupervised neural machine translation with weight sharing. *arXiv preprint arXiv:1804.09057*.
- Zhang, J.; Liu, Y.; Luan, H.; Xu, J.; and Sun, M. 2017. Prior knowledge integration for neural machine translation using posterior regularization. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 1514–1523.
- Zhang, Z.; Liu, S.; Li, M.; Zhou, M.; and Chen, E. 2018a. Joint training for neural machine translation models with monolingual data. In *AAAI*.
- Zhang, Z.; Wu, S.; Liu, S.; Li, M.; Zhou, M.; and Chen, E. 2018b. Regularizing neural machine translation by target-bidirectional agreement. *CoRR* abs/1808.04064.