

Bi-Directional Differentiable Input Reconstruction for Low-Resource Neural Machine Translation

Xing Niu

University of Maryland
xingniu@cs.umd.edu

Weijia Xu

University of Maryland
weijia@cs.umd.edu

Marine Carpuat

University of Maryland
marine@cs.umd.edu

Abstract

We aim to better exploit the limited amounts of parallel text available in low-resource settings by introducing a differentiable reconstruction loss for neural machine translation (NMT). This loss compares original inputs to reconstructed inputs, obtained by back-translating translation hypotheses into the input language. We leverage differentiable sampling and bi-directional NMT to train models end-to-end, without introducing additional parameters. This approach achieves small but consistent BLEU improvements on four language pairs in both translation directions, and outperforms an alternative differentiable reconstruction strategy based on hidden states.

1 Introduction

Neural Machine Translation (NMT) performance degrades sharply when parallel training data is limited (Koehn and Knowles, 2017). Past work has addressed this problem by leveraging monolingual data (Sennrich et al., 2016a; Ramachandran et al., 2017) or multilingual parallel data (Zoph et al., 2016; Johnson et al., 2017; Gu et al., 2018a). We hypothesize that the traditional training can be complemented by better leveraging limited training data. To this end, we propose a new training objective for this model by augmenting the standard translation cross-entropy loss with a **differentiable input reconstruction loss** to further exploit the source side of parallel samples.

Input reconstruction is motivated by the idea of round-trip translation. Suppose sentence f is translated forward to e using model θ_{fe} and then translated back to \hat{f} using model θ_{ef} , then e is more likely to be a good translation if the distance between \hat{f} and f is small (Brislin, 1970). Prior work applied round-trip translation to monolingual examples and sampled the intermediate translation e from a K -best list generated by model

θ_{fe} using beam search (Cheng et al., 2016; He et al., 2016). However, beam search is not differentiable which prevents back-propagating reconstruction errors to θ_{fe} . As a result, reinforcement learning algorithms, or independent updates to θ_{fe} and θ_{ef} were required.

In this paper, we focus on the problem of making input reconstruction differentiable to simplify training. In past work, Tu et al. (2017) addressed this issue by reconstructing source sentences from the decoder’s hidden states. However, this reconstruction task can be artificially easy if hidden states over-memorize the input. This approach also requires a separate auxiliary reconstructor, which introduces additional parameters.

We propose instead to combine benefits from differentiable sampling and bi-directional NMT to obtain a compact model that can be trained end-to-end with back-propagation. Specifically,

- Translations are sampled using the Straight-Through Gumbel Softmax (STGS) estimator (Jang et al., 2017; Bengio et al., 2013), which allows back-propagating reconstruction errors.
- Our approach builds on the bi-directional NMT model (Niu et al., 2018; Johnson et al., 2017), which improves low-resource translation by jointly modeling translation in both directions (e.g., Swahili \leftrightarrow English). A single bi-directional model is used as a translator and a reconstructor (i.e. $\theta_{ef} = \theta_{fe}$) without introducing more parameters.

Experiments show that our approach outperforms reconstruction from hidden states. It achieves consistent improvements across various low-resource language pairs and directions, showing its effectiveness in making better use of limited parallel data.

2 Background

Using round-trip translations ($f \rightarrow e \rightarrow \hat{f}$) as a training signal for NMT usually requires auxiliary models to perform back-translation and cannot be trained end-to-end without reinforcement learning. For instance, Cheng et al. (2016) added a reconstruction loss for monolingual examples to the training objective. He et al. (2016) evaluated the quality of e by a language model and \hat{f} by a reconstruction likelihood. Both approaches have symmetric forward and backward translation models which are updated alternatively. This requires policy gradient algorithms for training, which are not always stable.

Back-translation (Sennrich et al., 2016a) performs half of the reconstruction process, by generating a synthetic source side for monolingual target language examples: $e \rightarrow \hat{f}$. It uses an auxiliary backward model to generate the synthetic data but only updates the parameters of the primary forward model. Iteratively updating forward and backward models (Zhang et al., 2018; Niu et al., 2018) is an expensive solution as back-translations are regenerated at each iteration.

Prior work has sought to simplify the optimization of reconstruction losses by side-stepping beam search. Tu et al. (2017) first proposed to reconstruct NMT input from the decoder’s hidden states while Wang et al. (2018a,b) suggested to use both encoder and decoder hidden states to improve translation of dropped pronouns. However, these models might achieve low reconstruction errors by learning to copy the input to hidden states. To avoid copying the input, Artetxe et al. (2018) and Lample et al. (2018) used denoising autoencoders (Vincent et al., 2008) in unsupervised NMT.

Our approach is based instead on the Gumbel Softmax (Jang et al., 2017; Maddison et al., 2017), which facilitates differentiable sampling of sequences of discrete tokens. It has been successfully applied in many sequence generation tasks, including artificial language emergence for multi-agent communication (Havrylov and Titov, 2017), composing tree structures from text (Choi et al., 2018), and tasks under the umbrella of generative adversarial networks (Goodfellow et al., 2014) such as generating the context-free grammar (Kusner and Hernández-Lobato, 2016), machine comprehension (Wang et al., 2017) and machine translation (Gu et al., 2018b).

3 Approach

NMT is framed as a conditional language model, where the probability of predicting target token e_t at step t is conditioned on the previously generated sequence of tokens $e_{<t}$ and the source sequence f given the model parameter θ . Suppose each token is indexed and represented as a one-hot vector, its probability is realized as a softmax function over a linear transformation $a(\mathbf{h}_t)$ where \mathbf{h}_t is the decoder’s hidden state at step t :

$$P(e_t | e_{<t}, f; \theta) = \text{softmax}(a(\mathbf{h}_t))^\top e_t. \quad (1)$$

The hidden state is calculated by a neural network g given the embeddings of the previous target tokens $e_{<t}$ in the embedding matrix $E(e_{<t})$ and the context c_t coming from the source:

$$\mathbf{h}_t = g(E(e_{<t}), c_t). \quad (2)$$

In our bi-directional model, the source sentence can be either f or e and is respectively translated to e or f . The language is marked by a tag (e.g., $\langle \text{en} \rangle$) at the beginning of each source sentence (Johnson et al., 2017; Niu et al., 2018). To facilitate symmetric reconstruction, we also add language tags to target sentences. The training data corpus is then built by swapping the source and target sentences of a parallel corpus and appending the swapped version to the original.

3.1 Bi-Directional Reconstruction

Our bi-directional model performs both forward translation and backward reconstruction. By contrast, uni-directional models require an auxiliary reconstruction module, which introduces additional parameters. This module can be either a decoder-based reconstructor (Tu et al., 2017; Wang et al., 2018a,b) or a reversed dual NMT model (Cheng et al., 2016; He et al., 2016; Wang et al., 2018c; Zhang et al., 2018).

Here the reconstructor, which shares the same parameter with the translator $T(\cdot)$, can also be trained end-to-end by maximizing the log-likelihood of reconstructing f :

$$\mathcal{L}_R = \sum_f \log P(f | T(f; \theta); \theta), \quad (3)$$

Combining with the forward translation likelihood

$$\mathcal{L}_T = \sum_{(f||e)} \log P(e | f; \theta), \quad (4)$$

we use $\mathcal{L} = \mathcal{L}_T + \mathcal{L}_R$ as the final training objective for $\mathbf{f} \rightarrow \mathbf{e}$. The dual $\mathbf{e} \rightarrow \mathbf{f}$ model is trained simultaneously by swapping the language direction in bi-directional NMT.

Reconstruction is reliable only with a model that produces reasonable base translations. Following prior work (Tu et al., 2017; He et al., 2016; Cheng et al., 2016), we pre-train a base model with \mathcal{L}_T and fine-tune it with $\mathcal{L}_T + \mathcal{L}_R$.

3.2 Differentiable Sampling

We use differentiable sampling to side-step beam search and back-propagate error signals. We use the Gumbel-Max reparameterization trick (Maddison et al., 2014) to sample a translation token at each time step from the softmax distribution in Equation 1:

$$e_t = \text{one-hot} \left(\arg \max_k (a(\mathbf{h}_t)_k + G_k) \right) \quad (5)$$

where G_k is i.i.d. and drawn from $\text{Gumbel}(0, 1)$ ¹. We use scaled Gumbel with parameter β , i.e. $\text{Gumbel}(0, \beta)$, to control the randomness. The sampling becomes deterministic (which is equivalent to greedy search) as β approaches 0.

Since $\arg \max$ is not a differentiable operation, we approximate its gradient with the Straight-Through Gumbel Softmax (STGS) (Jang et al., 2017; Bengio et al., 2013): $\nabla_{\theta} e_t \approx \nabla_{\theta} \tilde{e}_t$, where

$$\tilde{e}_t = \text{softmax} \left((a(\mathbf{h}_t) + G) / \tau \right) \quad (6)$$

As τ approaches 0, softmax is closer to $\arg \max$ but training might be more unstable. While the STGS estimator is biased when τ is large, it performs well in practice (Gu et al., 2018b; Choi et al., 2018) and is sometimes faster and more effective than reinforcement learning (Havrylov and Titov, 2017).

To generate coherent intermediate translations, the decoder used for sampling only consumes its previously predicted $\hat{e}_{<t}$. This contrasts with the usual *teacher forcing* strategy (Williams and Zipser, 1989), which always feeds in the ground-truth previous tokens $e_{<t}$ when predicting the current token \hat{e}_t . With teacher forcing, the sequence concatenation $[e_{<t}; \hat{e}_t]$ is probably coherent at each time step, but the actual predicted sequence $[\hat{e}_{<t}; \hat{e}_t]$ would break the continuity.²

¹i.e. $G_k = -\log(-\log(u_k))$ and $u_k \sim \text{Uniform}(0, 1)$.

²Sampling with teacher forcing yielded consistently worse BLEU than baselines in preliminary experiments.

# sent.	Training	Dev.	Test
SW↔EN	60,570	500	3,000
TL↔EN	70,703	704	3,000
SO↔EN	68,550	844	3,000
TR↔EN	207,021	1,001	3,007

Table 1: Experiments are conducted on four low-resource language pairs, in both translation directions.

4 Experiments

4.1 Tasks and Data

We evaluate our approach on four low-resource language pairs. Parallel data for Swahili↔English (SW↔EN), Tagalog↔English (TL↔EN) and Somali↔English (SO↔EN) contains a mixture of domains such as news and weblogs and is collected from the IARPA MATERIAL program³, the Global Voices parallel corpus⁴, Common Crawl (Smith et al., 2013), and the LORELEI Somali representative language pack (LDC2018T11). The test samples are extracted from the held-out ANALYSIS set of MATERIAL. Parallel Turkish↔English (TR↔EN) data is provided by the WMT news translation task (Bojar et al., 2018). We use pre-processed “corpus”, “newsdev2016”, “newstest2017” as training, development and test sets.⁵

We apply normalization, tokenization, true-casing, joint source-target BPE with 32,000 operations (Sennrich et al., 2016b) and sentence-filtering (length 80 cutoff) to parallel data. Itemized data statistics after preprocessing can be found in Table 1. We report case-insensitive BLEU with the WMT standard ‘13a’ tokenization using SacreBLEU (Post, 2018).

4.2 Model Configuration and Baseline

We build NMT models upon the attentional RNN encoder-decoder architecture (Bahdanau et al., 2015) implemented in the Sockeye toolkit (Hieber et al., 2017). Our translation model uses a bi-directional encoder with a single LSTM layer of size 512, multilayer perceptron attention with a layer size of 512, and word representations of size 512. We apply layer normalization (Ba et al.,

³<https://www.iarpa.gov/index.php/research-programs/material>

⁴<http://casmacat.eu/corpus/global-voices.html>

⁵<http://data.statmt.org/wmt18/translation-task/preprocessed/>

Model	EN→SW	SW→EN	EN→TL	TL→EN	EN→SO	SO→EN	EN→TR	TR→EN
Baseline	33.60 ± 0.14	30.70 ± 0.19	27.23 ± 0.11	32.15 ± 0.21	12.25 ± 0.08	20.80 ± 0.12	12.90 ± 0.04	15.32 ± 0.11
HIDDEN	33.41 ± 0.15	30.91 ± 0.19	27.43 ± 0.14	32.20 ± 0.35	12.30 ± 0.11	20.72 ± 0.16	12.77 ± 0.11	15.34 ± 0.10
Δ	-0.19 ± 0.24	0.21 ± 0.14	0.19 ± 0.13	0.04 ± 0.17	0.05 ± 0.11	-0.08 ± 0.12	-0.13 ± 0.13	0.01 ± 0.07
β = 0	33.92 ± 0.10	31.37 ± 0.18	27.65 ± 0.09	32.75 ± 0.32	12.47 ± 0.08	21.14 ± 0.19	13.26 ± 0.07	15.60 ± 0.19
Δ	0.32 ± 0.12	0.66 ± 0.11	0.42 ± 0.16	0.59 ± 0.13	0.22 ± 0.04	0.35 ± 0.15	0.36 ± 0.09	0.28 ± 0.11
β = 0.5	33.97 ± 0.08	31.39 ± 0.09	27.65 ± 0.10	32.65 ± 0.24	12.48 ± 0.09	21.20 ± 0.14	13.16 ± 0.08	15.52 ± 0.07
Δ	0.37 ± 0.09	0.69 ± 0.11	0.42 ± 0.11	0.50 ± 0.08	0.23 ± 0.03	0.41 ± 0.13	0.25 ± 0.09	0.19 ± 0.05

Table 2: BLEU scores on eight translation directions. The numbers before and after ‘±’ are the mean and standard deviation over five randomly seeded models. Our proposed methods ($\beta = 0/0.5$) achieve small but consistent improvements. Δ BLEU scores are in bold if mean−std is above zero while in red if the mean is below zero.

2016) and add dropout to embeddings and RNNs (Gal and Ghahramani, 2016) with probability 0.2. We train using the Adam optimizer (Kingma and Ba, 2015) with a batch size of 48 sentences and we checkpoint the model every 1000 updates. The learning rate for baseline models is initialized to 0.001 and reduced by 30% after 4 checkpoints without improvement of perplexity on the development set. Training stops after 10 checkpoints without improvement.

The bi-directional NMT model ties source and target embeddings to yield a bilingual vector space. It also ties the output layer’s weights and embeddings to achieve better performance in low-resource scenarios (Press and Wolf, 2017; Nguyen and Chiang, 2018).

We train five randomly seeded bi-directional baseline models by optimizing the forward translation objective \mathcal{L}_T and report the mean and standard deviation of test BLEU. We fine-tune baseline models with objective $\mathcal{L}_T + \mathcal{L}_R$, inheriting all settings except the learning rate which is re-initialized to 0.0001. Each randomly seeded model is fine-tuned independently, so we are able to report the standard deviation of Δ BLEU.

4.3 Contrastive Reconstruction Model

We compare our approach with reconstruction from hidden states (HIDDEN). Following the best practice of Wang et al. (2018a), two reconstructors are used to take hidden states from both the encoder and the decoder. The corresponding two reconstruction losses and the canonical translation loss were originally uniformly weighted (i.e. 1, 1, 1), but we found that balancing the reconstruction and translation losses yields better results (i.e. 0.5, 0.5, 1) in preliminary experiments.⁶

We use the reconstructor exclusively to compute the reconstruction training loss. It has also been

⁶We observed around 0.2 BLEU gains for TR↔EN tasks.

used to re-rank translation hypotheses in prior work, but Tu et al. (2017) showed in ablation studies that the gains from re-ranking are small compared to those from training.

4.4 Results

Table 2 shows that our reconstruction approach achieves small but consistent BLEU improvements over the baseline on all eight tasks.⁷

We evaluate the impact of the Gumbel Softmax hyperparameters on the development set. We select $\tau = 2$ and $\beta = 0/0.5$ based on training stability and BLEU. Greedy search (i.e. $\beta = 0$) performs similarly as sampling with increased Gumbel noise (i.e. more random translation selection when $\beta = 0.5$): increased randomness in sampling does not have a strong impact on BLEU, even though random sampling may approximate the data distribution better (Ott et al., 2018). We hypothesize that more random translation selection introduces lower quality samples and therefore noisier training signals. This is consistent with the observation that random sampling is less effective for back-translation in low-resource settings (Edunov et al., 2018).

Sampling-based reconstruction is effective even if there is moderate domain mismatch between the training and the test data, such as in the case that the word type out-of-vocabulary (OOV) rate of TR→EN is larger than 20%. Larger improvements can be achieved when the test data is closer to training examples. For example, the OOV rate of SW→EN is much smaller than the OOV rate of TR→EN and the former obtains higher Δ BLEU.

Our approach yields more consistent results than reconstructing from hidden states. The latter fails to improve BLEU in more difficult cases, such as TR↔EN with high OOV rates. We observe extremely low training perplexity for HID-

⁷The improvements are significant with $p < 0.01$.




Figure 1: Training curves of perplexity on the training and the development sets for $TR \leftrightarrow EN$. Reconstructing from hidden states (HIDDEN) and reconstructing from sampled translations ($\beta = 0$) are compared. HIDDEN achieves extremely low training perplexity and suffers from unstable training during the early stage.

DEN compared with our proposed approach (Figure 1a). This suggests that HIDDEN yields representations that memorize the input rather than improve output representations.

Another advantage of our approach is that all parameters were jointly pre-trained, which results in more stable training behavior. By contrast, reconstructing from hidden states requires to initialize the reconstructors independently and suffers from unstable early training behavior (Figure 1).

5 Conclusion

We studied reconstructing the input of NMT from its intermediate translations to better exploit training samples in low-resource settings. We used a bi-directional NMT model and the Straight-Through Gumbel Softmax to build a fully differentiable reconstruction model that does not require any additional parameters. We empirically demonstrated that our approach is effective in low-resource scenarios. In future work, we will investigate the use of differentiable reconstruction from sampled sequences in unsupervised and semi-supervised sequence generation tasks. In particular, we will exploit monolingual corpora in addition to parallel corpora for NMT.

Acknowledgments

We thank the three anonymous reviewers for their helpful comments and suggestions. We also thank the members of the Computational Linguistics and Information Processing (CLIP) lab at the University of Maryland for helpful discussions.

This research is based upon work supported in part by an Amazon Web Services Machine Learning Research Award, and by the Office of

the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via contract #FA8650-17-C-9117. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- Mikel Artetxe, Gorika Labaka, Eneko Agirre, and Kyunghyun Cho. 2018. [Unsupervised neural machine translation](#). In *Proceedings of the 6th International Conference on Learning Representations*.
- Lei Jimmy Ba, Ryan Kiros, and Geoffrey E. Hinton. 2016. [Layer normalization](#). *CoRR*, abs/1607.06450.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *Proceedings of the 3th International Conference on Learning Representations*.
- Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. [Estimating or propagating gradients through stochastic neurons for conditional computation](#). *CoRR*, abs/1308.3432.
- Ondrej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, and Christof Monz. 2018. [Findings of the 2018 conference on machine translation \(WMT18\)](#). In *Proceedings of the Third Conference on Machine Translation*, pages 272–303. Association for Computational Linguistics.
- Richard W. Brislin. 1970. [Back-translation for cross-cultural research](#). *Journal of Cross-Cultural Psychology*, 1(3):185–216.

- Yong Cheng, Wei Xu, Zhongjun He, Wei He, Hua Wu, Maosong Sun, and Yang Liu. 2016. [Semi-supervised learning for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1965–1974. Association for Computational Linguistics.
- Jihun Choi, Kang Min Yoo, and Sang-goo Lee. 2018. [Learning to compose task-specific tree structures](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5094–5101. AAAI Press.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500. Association for Computational Linguistics.
- Yarin Gal and Zoubin Ghahramani. 2016. [A theoretically grounded application of dropout in recurrent neural networks](#). In *Advances in Neural Information Processing Systems 29*, pages 1019–1027. Curran Associates, Inc.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. 2014. [Generative adversarial nets](#). In *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc.
- Jiatao Gu, Hany Hassan, Jacob Devlin, and Victor O. K. Li. 2018a. [Universal neural machine translation for extremely low resource languages](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 344–354. Association for Computational Linguistics.
- Jiatao Gu, Daniel Jiwoong Im, and Victor O. K. Li. 2018b. [Neural machine translation with gumbel-greedy decoding](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5125–5132. AAAI Press.
- Serhii Havrylov and Ivan Titov. 2017. [Emergence of language with multi-agent games: Learning to communicate with sequences of symbols](#). In *Advances in Neural Information Processing Systems 30*, pages 2146–2156. Curran Associates, Inc.
- Di He, Yingce Xia, Tao Qin, Liwei Wang, Nenghai Yu, Tie-Yan Liu, and Wei-Ying Ma. 2016. [Dual learning for machine translation](#). In *Advances in Neural Information Processing Systems 29*, pages 820–828. Curran Associates, Inc.
- Felix Hieber, Tobias Domhan, Michael Denkowski, David Vilar, Artem Sokolov, Ann Clifton, and Matt Post. 2017. [Sockeye: A toolkit for neural machine translation](#). *CoRR*, abs/1712.05690.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. [Categorical reparameterization with gumbel-softmax](#). In *Proceedings of the 5th International Conference on Learning Representations*.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proceedings of the 3th International Conference on Learning Representations*.
- Philipp Koehn and Rebecca Knowles. 2017. [Six challenges for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 28–39. Association for Computational Linguistics.
- Matt J. Kusner and José Miguel Hernández-Lobato. 2016. [GANS for sequences of discrete elements with the gumbel-softmax distribution](#). In *Proceedings of the NIPS 2016 Workshop on Adversarial Training*.
- Guillaume Lample, Ludovic Denoyer, and Marc’Aurelio Ranzato. 2018. [Unsupervised machine translation using monolingual corpora only](#). In *Proceedings of the 6th International Conference on Learning Representations*.
- Chris J. Maddison, Andriy Mnih, and Yee Whye Teh. 2017. [The concrete distribution: A continuous relaxation of discrete random variables](#). In *Proceedings of the 5th International Conference on Learning Representations*.
- Chris J. Maddison, Daniel Tarlow, and Tom Minka. 2014. [A* sampling](#). In *Advances in Neural Information Processing Systems 27*, pages 3086–3094. Curran Associates, Inc.
- Toan Q. Nguyen and David Chiang. 2018. [Improving lexical choice in neural machine translation](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 334–343. Association for Computational Linguistics.
- Xing Niu, Michael Denkowski, and Marine Carpuat. 2018. [Bi-directional neural machine translation with synthetic parallel data](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 84–91. Association for Computational Linguistics.

- Myle Ott, Michael Auli, David Grangier, and Marc'Aurelio Ranzato. 2018. [Analyzing uncertainty in neural machine translation](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 3953–3962. PMLR.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation*, pages 186–191. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics*, pages 157–163. Association for Computational Linguistics.
- Prajit Ramachandran, Peter J. Liu, and Quoc V. Le. 2017. [Unsupervised pretraining for sequence to sequence learning](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 383–391. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 86–96. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1715–1725. Association for Computational Linguistics.
- Jason R. Smith, Herve Saint-Amand, Magdalena Plamada, Philipp Koehn, Chris Callison-Burch, and Adam Lopez. 2013. [Dirt cheap web-scale parallel text from the common crawl](#). In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, pages 1374–1383. Association for Computational Linguistics.
- Zhaopeng Tu, Yang Liu, Lifeng Shang, Xiaohua Liu, and Hang Li. 2017. [Neural machine translation with reconstruction](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 3097–3103. AAAI Press.
- Pascal Vincent, Hugo Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. 2008. [Extracting and composing robust features with denoising autoencoders](#). In *Proceedings of the 25th International Conference on Machine Learning*, pages 1096–1103. ACM.
- Bingning Wang, Kang Liu, and Jun Zhao. 2017. [Conditional generative adversarial networks for commonsense machine comprehension](#). In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, pages 4123–4129.
- Longyue Wang, Zhaopeng Tu, Shuming Shi, Tong Zhang, Yvette Graham, and Qun Liu. 2018a. [Translating pro-drop languages with reconstruction models](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 4937–4945. AAAI Press.
- Longyue Wang, Zhaopeng Tu, Andy Way, and Qun Liu. 2018b. [Learning to jointly translate and predict dropped pronouns with a shared reconstruction mechanism](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2997–3002. Association for Computational Linguistics.
- Yijun Wang, Yingce Xia, Li Zhao, Jiang Bian, Tao Qin, Guiquan Liu, and Tie-Yan Liu. 2018c. [Dual transfer learning for neural machine translation with marginal distribution regularization](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 5553–5560. AAAI Press.
- Ronald J. Williams and David Zipser. 1989. [A learning algorithm for continually running fully recurrent neural networks](#). *Neural Computation*, 1(2):270–280.
- Zhirui Zhang, Shujie Liu, Mu Li, Ming Zhou, and En-hong Chen. 2018. [Joint training for neural machine translation models with monolingual data](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, pages 555–562. AAAI Press.
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. 2016. [Transfer learning for low-resource neural machine translation](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1568–1575. Association for Computational Linguistics.