

# Sparse and Constrained Attention for Neural Machine Translation

<b>Chaitanya Malaviya*</b> Language Technologies Institute Carnegie Mellon University cmalaviy@cs.cmu.edu	<b>Pedro Ferreira</b> Instituto Superior Técnico Universidade de Lisboa Lisbon, Portugal pedro.miguel.ferreira@tecnico.ulisboa.pt	<b>André F. T. Martins</b> Unbabel & Instituto de Telecomunicações Lisbon, Portugal andre.martins@unbabel.com
--	---	---

## Abstract

In NMT, words are sometimes dropped from the source or generated repeatedly in the translation. We explore novel strategies to address the coverage problem that change only the attention transformation. Our approach allocates fertilities to source words, used to bound the attention each word can receive. We experiment with various sparse and constrained attention transformations and propose a new one, constrained sparsemax, shown to be differentiable and sparse. Empirical evaluation is provided in three languages pairs.

## 1 Introduction

Neural machine translation (NMT) emerged in the last few years as a very successful paradigm (Sutskever et al., 2014; Bahdanau et al., 2014; Gehring et al., 2017; Vaswani et al., 2017). While NMT is generally more fluent than previous statistical systems, adequacy is still a major concern (Koehn and Knowles, 2017): common mistakes include dropping source words and repeating words in the generated translation.

Previous work has attempted to mitigate this problem in various ways. Wu et al. (2016) incorporate coverage and length penalties during beam search—a simple yet limited solution, since it only affects the scores of translation hypotheses that are already in the beam. Other approaches involve architectural changes: providing coverage vectors to track the attention history (Mi et al., 2016; Tu et al., 2016), using gating architectures and adaptive attention to control the amount of source context provided (Tu et al., 2017a; Li and Zhu, 2017), or adding a reconstruction loss (Tu et al., 2017b). Feng et al. (2016) also use the notion of fertility

implicitly in their proposed model. Their fertility conditioned decoder uses a coverage vector and an extract gate which are incorporated in the decoding recurrent unit, increasing the number of parameters.

In this paper, we propose a different solution that does not change the overall architecture, but only the **attention transformation**. Namely, we replace the traditional softmax by other recently proposed transformations that either promote attention sparsity (Martins and Astudillo, 2016) or upper bound the amount of attention a word can receive (Martins and Kreutzer, 2017). The bounds are determined by the fertility values of the source words. While these transformations have given encouraging results in various NLP problems, they have never been applied to NMT, to the best of our knowledge. Furthermore, we combine these two ideas and propose a novel attention transformation, **constrained sparsemax**, which produces *both* sparse and bounded attention weights, yielding a compact and interpretable set of alignments. While being in-between soft and hard alignments (Figure 2), the constrained sparsemax transformation is end-to-end differentiable, hence amenable for training with gradient backpropagation.

To sum up, our contributions are as follows:<sup>1</sup>

- We formulate constrained sparsemax and derive efficient linear and sublinear-time algorithms for running forward and backward propagation. This transformation has two levels of sparsity: over time steps, and over the attended words at each step.
- We provide a detailed empirical comparison of various attention transformations, including softmax (Bahdanau et al., 2014), sparse-

<sup>1</sup>Our software code is available at the OpenNMT fork [www.github.com/Unbabel/OpenNMT-py/tree/dev](http://www.github.com/Unbabel/OpenNMT-py/tree/dev) and the running scripts at [www.github.com/Unbabel/sparse\\_constrained\\_attention](http://www.github.com/Unbabel/sparse_constrained_attention).

\*Work done during an internship at Unbabel.

max (Martins and Astudillo, 2016), constrained softmax (Martins and Kreutzer, 2017), and our newly proposed constrained sparsemax. We provide error analysis including two new metrics targeted at detecting coverage problems.

## 2 Preliminaries

Our underlying model architecture is a standard attentional encoder-decoder (Bahdanau et al., 2014). Let  $x := x_{1:J}$  and  $y := y_{1:T}$  denote the source and target sentences, respectively. We use a Bi-LSTM encoder to represent the source words as a matrix  $\mathbf{H} := [\mathbf{h}_1, \dots, \mathbf{h}_J] \in \mathbb{R}^{2D \times J}$ . The conditional probability of the target sentence is given as

$$p(y | x) := \prod_{t=1}^T p(y_t | y_{1:(t-1)}, x), \quad (1)$$

where  $p(y_t | y_{1:(t-1)}, x)$  is computed by a softmax output layer that receives a decoder state  $\mathbf{s}_t$  as input. This state is updated by an auto-regressive LSTM,  $\mathbf{s}_t = \text{RNN}(\text{embed}(y_{t-1}), \mathbf{s}_{t-1}, \mathbf{c}_t)$ , where  $\mathbf{c}_t$  is an input context vector. This vector is computed as  $\mathbf{c}_t := \mathbf{H}\boldsymbol{\alpha}_t$ , where  $\boldsymbol{\alpha}_t$  is a probability distribution that represents the attention over the source words, commonly obtained as

$$\boldsymbol{\alpha}_t = \text{softmax}(\mathbf{z}_t), \quad (2)$$

where  $\mathbf{z}_t \in \mathbb{R}^J$  is a vector of scores. We follow Luong et al. (2015) and define  $z_{t,j} := \mathbf{s}_{t-1}^\top \mathbf{W}\mathbf{h}_j$  as a bilinear transformation of encoder and decoder states, where  $\mathbf{W}$  is a model parameter.<sup>2</sup>

## 3 Sparse and Constrained Attention

In this work, we consider alternatives to Eq. 2. Since the softmax is strictly positive, it forces all words in the source to receive *some* probability mass in the resulting attention distribution, which can be wasteful. Moreover, it may happen that the decoder attends repeatedly to the same source words across time steps, causing repetitions in the generated translation, as Tu et al. (2016) observed.

With this in mind, we replace Eq. 2 by  $\boldsymbol{\alpha}_t = \rho(\mathbf{z}_t, \mathbf{u}_t)$ , where  $\rho$  is a transformation that may depend both on the scores  $\mathbf{z}_t \in \mathbb{R}^J$  and on **upper bounds**  $\mathbf{u}_t \in \mathbb{R}^J$  that limit the amount of attention that each word can receive. We consider three alternatives to softmax, described next.

<sup>2</sup>This is the default implementation in the OpenNMT package. In preliminary experiments, feedforward attention (Bahdanau et al., 2014) did not show improvements.

**Sparsemax.** The sparsemax transformation (Martins and Astudillo, 2016) is defined as:

$$\text{sparsemax}(\mathbf{z}) := \arg \min_{\boldsymbol{\alpha} \in \Delta^J} \|\boldsymbol{\alpha} - \mathbf{z}\|^2, \quad (3)$$

where  $\Delta^J := \{\boldsymbol{\alpha} \in \mathbb{R}^J \mid \boldsymbol{\alpha} \geq \mathbf{0}, \sum_j \alpha_j = 1\}$ . In words, it is the Euclidean projection of the scores  $\mathbf{z}$  onto the probability simplex. These projections tend to hit the boundary of the simplex, yielding a sparse probability distribution. This allows the decoder to attend only to a few words in the source, assigning zero probability mass to all other words. Martins and Astudillo (2016) have shown that the sparsemax can be evaluated in  $O(J)$  time (same asymptotic cost as softmax) and gradient backpropagation takes sublinear time (faster than softmax), by exploiting the sparsity of the solution.

**Constrained softmax.** The constrained softmax transformation was recently proposed by Martins and Kreutzer (2017) in the context of easy-first sequence tagging, being defined as follows:

$$\begin{aligned} \text{csoftmax}(\mathbf{z}; \mathbf{u}) &:= \arg \min_{\boldsymbol{\alpha} \in \Delta^J} \text{KL}(\boldsymbol{\alpha} \parallel \text{softmax}(\mathbf{z})) \\ &\text{s.t. } \boldsymbol{\alpha} \leq \mathbf{u}, \end{aligned} \quad (4)$$

where  $\mathbf{u}$  is a vector of upper bounds, and  $\text{KL}(\cdot \parallel \cdot)$  is the Kullback-Leibler divergence. In other words, it returns the distribution closest to  $\text{softmax}(\mathbf{z})$  whose attention probabilities are bounded by  $\mathbf{u}$ . Martins and Kreutzer (2017) have shown that this transformation can be evaluated in  $O(J \log J)$  time and its gradients backpropagated in  $O(J)$  time.

To use this transformation in the attention mechanism, we make use of the idea of **fertility** (Brown et al., 1993). Namely, let  $\boldsymbol{\beta}_{t-1} := \sum_{\tau=1}^{t-1} \boldsymbol{\alpha}_\tau$  denote the **cumulative attention** that each source word has received up to time step  $t$ , and let  $\mathbf{f} := (f_j)_{j=1}^J$  be a vector containing fertility upper bounds for each source word. The attention at step  $t$  is computed as

$$\boldsymbol{\alpha}_t = \text{csoftmax}(\mathbf{z}_t, \mathbf{f} - \boldsymbol{\beta}_{t-1}). \quad (5)$$

Intuitively, each source word  $j$  gets a credit of  $f_j$  units of attention, which are consumed along the decoding process. If all the credit is exhausted, it receives zero attention from then on. Unlike the sparsemax transformation, which places sparse attention over the source words, the constrained softmax leads to sparsity over time steps.

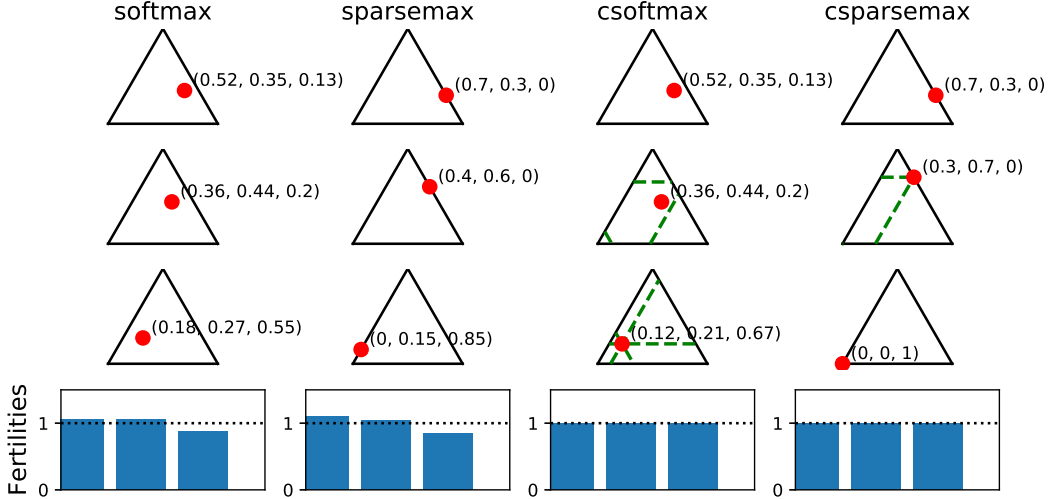


Figure 1: Illustration of the different attention transformations for a toy example with three source words. We show the attention values on the probability simplex. In the first row we assume scores  $z = (1.2, 0.8, -0.2)$ , and in the second and third rows  $z = (0.7, 0.9, 0.1)$  and  $z = (-0.2, 0.2, 0.9)$ , respectively. For constrained softmax/sparsemax, we set unit fertilities to every word; for each row the upper bounds (represented as green dashed lines) are set as the difference between these fertilities and the cumulative attention each word has received. The last row illustrates the cumulative attention for the three words after all rounds.

**Constrained sparsemax.** In this work, we propose a novel transformation which shares the two properties above: it provides *both* sparse and bounded probabilities. It is defined as:

$$\text{csparsemax}(z; \mathbf{u}) := \arg \min_{\alpha \in \Delta^J} \|\alpha - z\|^2 \quad \text{s.t. } \alpha \leq \mathbf{u}. \quad (6)$$

The following result, whose detailed proof we include as supplementary material (Appendix A), is key for enabling the use of the constrained sparsemax transformation in neural networks.

**Proposition 1** *Let  $\alpha^* = \text{csparsemax}(z; \mathbf{u})$  be the solution of Eq. 6, and define the sets  $\mathcal{A} = \{j \in [J] \mid 0 < \alpha_j^* < u_j\}$ ,  $\mathcal{A}_L = \{j \in [J] \mid \alpha_j^* = 0\}$ , and  $\mathcal{A}_R = \{j \in [J] \mid \alpha_j^* = u_j\}$ . Then:*

- **Forward propagation.**  $\alpha^*$  can be computed in  $O(J)$  time with the algorithm of [Pardalos and Koor \(1990\)](#) (Alg. 1 in Appendix A). The solution takes the form  $\alpha_j^* = \max\{0, \min\{u_j, z_j - \tau\}\}$ , where  $\tau$  is a normalization constant.
- **Gradient backpropagation.** Backpropagation takes sublinear time  $O(|\mathcal{A}| + |\mathcal{A}_R|)$ . Let  $L(\theta)$

be a loss function,  $d\alpha = \nabla_{\alpha} L(\theta)$  be the output gradient, and  $dz = \nabla_z L(\theta)$  and  $du = \nabla_u L(\theta)$  be the input gradients. Then, we have:

$$dz_j = \mathbf{1}(j \in \mathcal{A})(d\alpha_j - m) \quad (7)$$

$$du_j = \mathbf{1}(j \in \mathcal{A}_R)(d\alpha_j - m), \quad (8)$$

where  $m = \frac{1}{|\mathcal{A}|} \sum_{j \in \mathcal{A}} d\alpha_j$ .

## 4 Fertility Bounds

We experiment with three ways of setting the fertility of the source words: CONSTANT, GUIDED, and PREDICTED. With CONSTANT, we set the fertilities of all source words to a fixed integer value  $f$ . With GUIDED, we train a word aligner based on IBM Model 2 (we used `fast_align` in our experiments, [Dyer et al. \(2013\)](#)) and, for each word in the vocabulary, we set the fertilities to the maximal observed value in the training data (or 1 if no alignment was observed). With the PREDICTED strategy, we train a separate fertility predictor model using a bi-LSTM tagger.<sup>3</sup> At training time, we provide as supervision the fertility estimated by `fast_align`. Since our model works

<sup>3</sup>A similar strategy was recently used by [Gu et al. \(2018\)](#) as a component of their non-autoregressive NMT model.

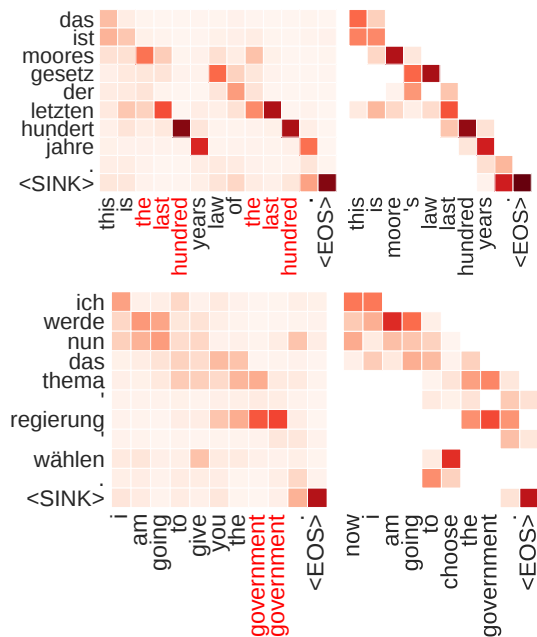


Figure 2: Attention maps for softmax and csparsemax for two DE-EN sentence pairs (white means zero attention). Repeated words are **high-lighted**. The reference translations are “*This is Moore’s law over the last hundred years*” and “*I am going to go ahead and select government.*”

with fertility *upper bounds* and the word aligner may miss some word pairs, we found it beneficial to add a constant to this number (1 in our experiments). At test time, we use the expected fertilities according to our model.

**Sink token.** We append an additional `<SINK>` token to the end of the source sentence, to which we assign unbounded fertility ( $f_{J+1} = \infty$ ). The token is akin to the null alignment in IBM models. The reason we add this token is the following: without the sink token, the length of the generated target sentence can never exceed  $\sum_j f_j$  words if we use constrained softmax/sparsemax. At training time this may be problematic, since the target length is fixed and the problems in Eqs. 4–6 can become infeasible. By adding the sink token we guarantee  $\sum_j f_j = \infty$ , eliminating the problem.

**Exhaustion strategies.** To avoid missing source words, we implemented a simple strategy to encourage more attention to words with larger credit: we redefine the pre-attention word scores as  $z'_t = z_t + c u_t$ , where  $c$  is a constant ( $c = 0.2$  in our experiments). This increases the score of words which have not yet exhausted their fertility (we may regard it as a “soft” lower bound in Eqs. 4–6).

## 5 Experiments

We evaluated our attention transformations on three language pairs. We focused on small datasets, as they are the most affected by coverage mistakes. We use the IWSLT 2014 corpus for DE-EN, the KFTT corpus for JA-EN (Neubig, 2011), and the WMT 2016 dataset for RO-EN. The training sets have 153,326, 329,882, and 560,767 parallel sentences, respectively. Our reason to prefer smaller datasets is that this regime is what brings more adequacy issues and demands more structural biases, hence it is a good test bed for our methods. We tokenized the data using the Moses scripts and preprocessed it with subword units (Sennrich et al., 2016) with a joint vocabulary and 32k merge operations. Our implementation was done on a fork of the OpenNMT-py toolkit (Klein et al., 2017) with the default parameters<sup>4</sup>. We used a validation set to tune hyperparameters introduced by our model. Even though our attention implementations are CPU-based using NumPy (unlike the rest of the computation which is done on the GPU), we did not observe any noticeable slowdown using multiple devices.

As baselines, we use softmax attention, as well as two recently proposed coverage models:

- **COVPENALTY** (Wu et al., 2016, §7). At test time, the hypotheses in the beam are rescored with a global score that includes a length and a coverage penalty.<sup>5</sup> We tuned  $\alpha$  and  $\beta$  with grid search on  $\{0.2k\}_{k=0}^5$ , as in Wu et al. (2016).
- **COVVECTOR** (Tu et al., 2016). At training and test time, coverage vectors  $\beta$  and additional parameters  $v$  are used to condition the next attention step. We adapted this to our bilinear attention by defining  $z_{t,j} = s_{t-1}^\top (\mathbf{W} \mathbf{h}_j + \mathbf{v} \beta_{t-1,j})$ . We also experimented combining the strategies above with the sparsemax transformation.

As evaluation metrics, we report tokenized BLEU, METEOR (Denkowski and Lavie (2014), as well as two new metrics that we describe next to account for over and under-translation.<sup>6</sup>

<sup>4</sup>We used a 2-layer LSTM, embedding and hidden size of 500, dropout 0.3, and the SGD optimizer for 13 epochs.

<sup>5</sup>Since our sparse attention can become 0 for some words, we extended the original coverage penalty by adding another parameter  $\epsilon$ , set to 0.1:  $\text{cp}(x; y) := \beta \sum_{j=1}^J \log \max\{\epsilon, \min\{1, \sum_{t=1}^{|y|} \alpha_{jt}\}\}$ .

<sup>6</sup>Both evaluation metrics are included in our software package at [www.github.com/Unbabel/sparse\\_constrained\\_attention](http://www.github.com/Unbabel/sparse_constrained_attention).

	De-En				Ja-En				Ro-En			
	BLEU	METEOR	REP	DROP	BLEU	METEOR	REP	DROP	BLEU	METEOR	REP	DROP
softmax	29.51	31.43	3.37	5.89	20.36	23.83	13.48	23.30	29.67	32.05	2.45	5.59
softmax + COVPENALTY	29.69	31.53	3.47	5.74	20.70	24.12	14.12	22.79	29.81	32.15	2.48	5.49
softmax + COVVECTOR	29.63	31.54	2.93	5.65	21.53	24.50	11.07	22.18	<b>30.08</b>	<b>32.22</b>	2.42	5.47
sparsemax	29.73	31.54	3.18	5.90	21.28	24.25	13.09	22.40	29.97	32.12	2.19	5.60
sparsemax + COVPENALTY	29.83	31.60	3.24	5.79	<b>21.64</b>	24.49	13.36	21.91	30.07	32.20	2.20	5.47
sparsemax + COVVECTOR	29.22	31.18	3.13	6.15	21.35	<b>24.74</b>	<b>10.11</b>	<b>21.25</b>	29.30	31.84	2.18	5.87
csoftmax ( $c = 0.2$ )	29.39	31.33	3.29	5.86	20.71	24.00	12.38	22.73	29.39	31.83	2.37	5.64
csparsemax ( $c = 0.2$ )	<b>29.85</b>	<b>31.76</b>	<b>2.67</b>	<b>5.23</b>	21.31	24.51	11.40	21.59	29.77	32.10	<b>1.98</b>	<b>5.44</b>

Table 1: BLEU, METEOR, REP and DROP scores on the test sets for different attention transformations.

	BLEU	METEOR
CONSTANT, $f = 2$	29.66	31.60
CONSTANT, $f = 3$	29.64	31.56
GUIDED,	29.56	31.45
PREDICTED, $c = 0$	29.78	31.60
PREDICTED, $c = 0.2$	<b>29.85</b>	<b>31.76</b>

Table 2: Impact of various fertility strategies for the csparsemax attention model (DE-EN).

**REP-score:** a new metric to count repetitions. Formally, given an  $n$ -gram  $s \in V^n$ , let  $t(s)$  and  $r(s)$  be the its frequency in the model translation and reference. We first compute a sentence-level score

$$\sigma(t, r) = \lambda_1 \sum_{s \in V^n, t(s) \geq 2} \max\{0, t(s) - r(s)\} + \lambda_2 \sum_{w \in V} \max\{0, t(w) - r(w)\}.$$

The REP-score is then given by summing  $\sigma(t, r)$  over sentences, normalizing by the number of words on the reference corpus, and multiplying by 100. We used  $n = 2$ ,  $\lambda_1 = 1$  and  $\lambda_2 = 2$ .

**DROP-score:** a new metric that accounts for possibly dropped words. To compute it, we first compute two sets of word alignments: from source to reference translation, and from source to the predicted translation. In our experiments, the alignments were obtained with `fast_align` (Dyer et al., 2013), trained on the training partition of the data. Then, the DROP-score computes the percentage of source words that aligned with some word from the reference translation, but not with any word from the predicted translation.

Table 1 shows the results. We can see that on average, the sparse models (csparsemax as well as sparsemax combined with coverage models) have higher scores on both BLEU and METEOR. Generally, they also obtain better REP and DROP scores than csoftmax and softmax, which suggests that sparse attention alleviates the problem of coverage to some extent.

To compare different fertility strategies, we ran experiments on the DE-EN for the csparsemax transformation (Table 2). We see that the PREDICTED strategy outperforms the others both in terms of BLEU and METEOR, albeit slightly.

Figure 2 shows examples of sentences for which the csparsemax fixed repetitions, along with the corresponding attention maps. We see that in the case of softmax repetitions, the decoder attends repeatedly to the same portion of the source sentence (the expression “*letzten hundert*” in the first sentence and “*regierung*” in the second sentence). Not only did csparsemax avoid repetitions, but it also yielded a sparse set of alignments, as expected. Appendix B provides more examples of translations from all models in discussion.

## 6 Conclusions

We proposed a new approach to address the coverage problem in NMT, by replacing the softmax attentional transformation by sparse and constrained alternatives: sparsemax, constrained softmax, and the newly proposed constrained sparsemax. For the latter, we derived efficient forward and backward propagation algorithms. By incorporating a model for fertility prediction, our attention transformations led to sparse alignments, avoiding repeated words in the translation.

## Acknowledgments

We thank the Unbabel AI Research team for numerous discussions, and the three anonymous reviewers for their insightful comments. This work was supported by the European Research Council (ERC StG DeepSPIN 758969) and by the Fundação para a Ciência e Tecnologia through contracts UID/EEA/50008/2013, PTDC/EEI-SII/7092/2014 (LearnBig), and CMUPERI/TIC/0046/2014 (GoLocal).

## References

- Miguel B. Almeida and André F. T. Martins. 2013. Fast and Robust Compressive Summarization with Dual Decomposition and Multi-Task Learning. In *Proc. of the Annual Meeting of the Association for Computational Linguistics*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Manuel Blum, Robert W Floyd, Vaughan Pratt, Ronald L Rivest, and Robert E Tarjan. 1973. Time bounds for selection. *Journal of Computer and System Sciences* 7(4):448–461.
- Peter F Brown, Vincent J Della Pietra, Stephen A Della Pietra, and Robert L Mercer. 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational linguistics* 19(2):263–311.
- Michael Denkowski and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the EACL 2014 Workshop on Statistical Machine Translation*.
- Chris Dyer, Victor Chahuneau, and Noah A Smith. 2013. A simple, fast, and effective reparameterization of ibm model 2. Association for Computational Linguistics.
- Shi Feng, Shujie Liu, Nan Yang, Mu Li, Ming Zhou, and Kenny Q. Zhu. 2016. Improving attention modeling with implicit distortion and fertility for machine translation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. The COLING 2016 Organizing Committee, Osaka, Japan, pages 3082–3092.
- Jonas Gehring, Michael Auli, David Grangier, and Yann Dauphin. 2017. A convolutional encoder model for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. volume 1, pages 123–135.
- Jiatao Gu, James Bradbury, Caiming Xiong, Victor OK Li, and Richard Socher. 2018. Non-autoregressive neural machine translation. In *Proc. of International Conference on Learning Representations*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M. Rush. 2017. Opennmt: Open-source toolkit for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, System Demonstrations*. pages 67–72.
- Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*. Association for Computational Linguistics, Vancouver, pages 28–39.
- Junhui Li and Muhua Zhu. 2017. Learning when to attend for neural machine translation. *arXiv preprint arXiv:1705.11160*.
- Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. pages 1412–1421.
- Andre Martins and Ramon Astudillo. 2016. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *International Conference on Machine Learning*. pages 1614–1623.
- André FT Martins and Julia Kreutzer. 2017. Learning what’s easy: Fully differentiable neural easy-first taggers. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. pages 349–362.
- Haitao Mi, Baskaran Sankaran, Zhiguo Wang, and Abe Ittycheriah. 2016. Coverage embedding models for neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*. pages 955–960.
- Graham Neubig. 2011. The Kyoto free translation task. <http://www.phontron.com/kftt>.
- Panos M. Pardalos and Naina Kover. 1990. An algorithm for a singly constrained class of quadratic programs subject to upper and lower bounds. *Mathematical Programming* 46(1):321–328.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*. pages 3104–3112.
- Zhaopeng Tu, Yang Liu, Zhengdong Lu, Xiaohua Liu, and Hang Li. 2017a. Context gates for neural machine translation.
- Zhaopeng Tu, Yang Liu, Lifeng Shang, Xiaohua Liu, and Hang Li. 2017b. Neural machine translation with reconstruction. In *AAAI*. pages 3097–3103.
- Zhaopeng Tu, Zhengdong Lu, Yang Liu, Xiaohua Liu, and Hang Li. 2016. Modeling coverage for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. pages 6000–6010.

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.