

# Extreme Adaptation for Personalized Neural Machine Translation

**Paul Michel**

Language Technologies Institute  
Carnegie Mellon University  
pmichel1@cs.cmu.edu

**Graham Neubig**

Language Technologies Institute  
Carnegie Mellon University  
gneubig@cs.cmu.edu

## Abstract

Every person speaks or writes their own flavor of their native language, influenced by a number of factors: the content they tend to talk about, their gender, their social status, or their geographical origin. When attempting to perform Machine Translation (MT), these variations have a significant effect on how the system should perform translation, but this is not captured well by standard one-size-fits-all models. In this paper, we propose a simple and parameter-efficient adaptation technique that only requires adapting the bias of the output softmax to each particular user of the MT system, either directly or through a factored approximation. Experiments on TED talks in three languages demonstrate improvements in translation accuracy, and better reflection of speaker traits in the target text.

## 1 Introduction

The production of language varies depending on the speaker or author, be it to reflect personal traits (*e.g.* job, gender, role, dialect) or the topics that tend to be discussed (*e.g.* technology, law, religion). Current Neural Machine Translation (NMT) systems do not incorporate any explicit information about the speaker, and this forces the model to learn these traits implicitly. This is a difficult and indirect way to capture inter-personal variations, and in some cases it is impossible without external context (Table 1, Mirkin et al. (2015)).

Recent work has incorporated side information about the author such as personality (Mirkin et al., 2015), gender (Rabinovich et al., 2017) or politeness (Sennrich et al., 2016a), but these methods can only handle phenomena where there are ex-

Source	Translation
I went home	[Man]: Je suis <b>rentré</b> à la maison [Woman]: Je suis <b>rentrée</b> à la maison
I do drug testing	[Doctor]: Je <b>teste des médicaments</b> [Police]: Je <b>dépiste des drogues</b>

Table 1: Examples where speaker information influences English-French translation.

PLICIT labels for the traits. Our work investigates how we can efficiently model speaker-related variations to improve NMT models.

In particular, we are interested in improving our NMT system given few training examples for any particular speaker. We propose to approach this task as a domain adaptation problem with an extremely large number of domains and little data for each domain, a setting where we may expect traditional approaches to domain adaptation that adjust all model parameters to be sub-optimal (§2). Our proposed solution involves modeling the speaker-specific variations as an additional bias vector in the softmax layer, where we either learn this bias directly, or through a factored model that treats each user as a mixture of a few prototypical bias vectors (§3).

We construct a new dataset of Speaker Annotated TED talks (SATED, §4) to validate our approach. Adaptation experiments (§5) show that explicitly incorporating speaker information into the model improves translation quality and accuracy with respect to speaker traits.<sup>1</sup>

## 2 Problem Formulation and Baselines

In the rest of this paper, we refer to the person producing the source sentence (speaker, author,

<sup>1</sup>Data/code publicly available at <http://www.cs.cmu.edu/~pmichel1/sated/> and <https://github.com/neulab/extreme-adaptation-for-personalized-translation> respectively.

etc...) generically as the *speaker*. We denote as  $\mathcal{S}$  the set of all speakers.

The usual objective of NMT is to find parameters  $\theta$  of the conditional distribution  $p(y | x; \theta)$  to maximize the empirical likelihood. We argue that personal variations in language warrant decomposing the empirical distribution into  $|\mathcal{S}|$  speaker specific domains  $\mathcal{D}_s$  and learning a different set of parameters  $\theta_s$  for each. This setting exhibits specific traits that set it apart from common domain adaptation settings:

1. The number of speakers is very large. Our particular setting deals with  $|\mathcal{S}| \approx 1800$  but our approaches should be able to accommodate orders of magnitude more speakers.
2. There is very little data (even monolingual, let alone bilingual or parallel) for each speaker, compared to millions of sentences usually used in NMT.
3. As a consequence of 1, we can assume that many speakers share similar characteristics such as gender, social status, and as such may have similar associated domains.<sup>2</sup>

## 2.1 Baseline NMT model

All of our experiments are based on a standard neural sequence to sequence model. We use one layer LSTMs as the encoder and decoder and the *concat* attention mechanism described in Luong and Manning (2015). We share the parameters in the embedding and softmax matrix of the decoder as proposed in Press and Wolf (2017). All the layers have dimension 512 except for the attention layer (dimension 256). To make our baseline competitive, we apply several regularization techniques such as dropout (Srivastava et al., 2014) in the output layer and within the LSTM (using the variant presented in Gal and Ghahramani, 2016). We also drop words in the target sentence with probability 0.1 according to Iyyer et al. (2015) and implement label smoothing as proposed in Szegedy et al. (2016) with coefficient 0.1. Appendix A provides a more thorough description of the baseline model.

## 2.2 Baseline adaptation strategy

As mentioned in §2, our goal is to learn a separate conditional distribution  $p(y | x, s)$  and

<sup>2</sup>Note that the speakers are still unique, and many might use very specific words (e.g. the name of their company or of a specific medical procedure that they are an expert on).

parametrization  $\theta_s$  to improve translation for speaker  $s$ . The usual way of adapting from general domain parameters  $\theta$  to  $\theta_s$  is to retrain the full model on the domain specific data (Luong and Manning, 2015). Naively applying this approach in the context of personalizing a model for each speaker however has two main drawbacks:

**Parameter cost** Maintaining a set of model parameters for each speaker is expensive. For example, the model in §2.1 has  $\approx 47$ M parameters when the vocabulary size is 40k, as is the case in our experiments in §5. Assuming each parameter is stored as a 32bit float, every speaker-specific model costs  $\approx 188$ MB. In a production environment with thousands to billions of speakers, this is impractical.

**Overfitting** Training each speaker model with very little data is a challenge, necessitating careful and heavy regularization (Miceli Barone et al., 2017) and an early stopping procedure.

## 2.3 Domain Token

A more efficient domain adaptation technique is the *domain token* idea used in Sennrich et al. (2016a); Chu et al. (2017): introduce an additional token marking the domain in the source and/or the target sentence. In experiments, we add a token indicating the speaker at the start of the target sentence for each speaker. We refer to this method as the *spk\_token* method in the following.

Note that in this case there is now only an embedding vector (of dimension 512 in our experiments) for each speaker. However, the resulting domain embedding are non-trivial to interpret (i.e. it is not clear what they tell us about the domain or speaker itself).

## 3 Speaker-specific Vocabulary Bias

In NMT models, the final choice of which word to use in the next step  $t$  of translation is generally performed by the following softmax equation

$$p_t = \text{softmax}(E_T o_t + b_T) \quad (1)$$

where  $o_t$  is predicted in a context-sensitive manner by the NMT system and  $E_T$  and  $b_T$  are the weight matrix and bias vector parameters respectively. Importantly,  $b_T$  governs the overall likelihood that the NMT model will choose particular vocabulary. In this section, we describe our proposed methods for making this bias term speaker-

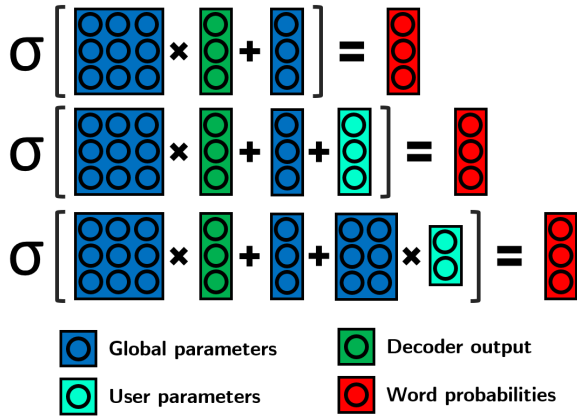


Figure 1: Graphical representation of our different adaptation models for the softmax layer. From top to bottom is the base softmax, the full\_bias softmax and the fact\_bias softmax

specific, which provides an efficient way to allow for speaker-specific vocabulary choice.<sup>3</sup>

### 3.1 Full speaker bias

We first propose to learn speaker-specific parameters for the bias term in the output softmax only. This means changing Eq. 1 to

$$p_t = \text{softmax}(E_T o_t + b_T + b_s) \quad (2)$$

for speaker  $s$ . This only requires learning and storing a vector equal to the size of the vocabulary, which is a mere 0.09% of the parameters in the full model in our experiments. In effect, this greatly reducing the parameter cost and concerns of overfitting cited in §2.2. This model is also easy to interpret as each coordinate of the bias vector corresponds to a log-probability on the target vocabulary. We refer to this variant as full\_bias.

### 3.2 Factored speaker bias

The biases for a set of speakers  $\mathcal{S}$  on a vocabulary  $\mathcal{V}$  can be represented as a matrix:

$$B \in \mathbb{R}^{|\mathcal{S}| \times |\mathcal{V}|} \quad (3)$$

where each row of  $B$  is one speaker bias  $b_s$ . In this formulation, the  $|\mathcal{S}|$  rows are still linearly independent, meaning that  $B$  is high rank. In practical terms, this means that we cannot share information among users about how their vocabulary

<sup>3</sup>Notably, while this limits the model to only handling word choice and does not explicitly allow it to model syntactic variations, favoring certain words over others can indirectly favor certain phenomena (e.g. favoring passive speech by increasing the probability of auxiliaries).

	en-fr	en-es	en-de
#talks	1,887	1,922	1,670
#train	177,743	182,582	156,134
#dev	3,774	3,844	3,340
#test	3,774	3,844	3,340
avg. sent/talk	94,2	95.0	93,5
std dev	57,6	57.8	60,3

Table 2: Dataset statistics

selection co-varies, which is likely sub-ideal given that speakers share common characteristics.

Thus, we propose another parametrization of the speaker bias, fact\_bias, where the  $B$  matrix is factored according to:

$$\begin{aligned} B &= S\tilde{B} \\ S &\in \mathbb{R}^{|\mathcal{S}| \times r}, \\ \tilde{B} &\in \mathbb{R}^{r \times |\mathcal{V}|} \end{aligned} \quad (4)$$

where  $S$  is a matrix of speaker vectors of low dimension  $r$  and  $\tilde{B}$  is a matrix of  $r$  speaker independent biases. Here, the bias for each speaker is a mixture of  $r$  ‘‘centroid’’ biases  $\tilde{B}$  with  $r$  speaker ‘‘weights’’. This reduces the total number of parameters allocated to speaker adaptation from  $|\mathcal{S}||\mathcal{V}|$  to  $r(|\mathcal{S}| + |\mathcal{V}|)$ . In our experiments, this corresponds to using between 99.38 and 99.45% fewer parameters than the full\_bias model depending on the language pair, with  $r$  parameters per speaker. In this work, we will use  $r = 10$ .

We provide a graphical summary of our proposed approaches in figure 1.

## 4 Speaker Annotated TED Talks Dataset

In order to evaluate the effectiveness of our proposed methods, we construct a new dataset, Speaker Annotated TED (SATED) based on TED talks,<sup>4</sup> with three language pairs, English-French (en-fr), English-German (en-de) and English-Spanish (en-es) and speaker annotation.

The dataset consists of transcripts directly collected from <https://www.ted.com/talks>, and contains roughly 271K sentences in each language distributed among 2324 talks. We pre-process the data by removing sentences that don’t have any translation or are longer than 60 words, lower-casing, and tokenizing (using the Moses tokenizer (Koehn et al., 2007)).

<sup>4</sup><https://www.ted.com>

Some talks are partially or not translated in some of the languages (in particular there are fewer translations in German than in French or Spanish), we therefore remove any talk with less than 10 translated sentences in each language pair.

The data is then partitioned into training, validation and test sets. We split the corpus such that the test and validation split each contain 2 sentence pairs from each talk, thus ensuring that all talks are present in every split. Each sentence pair is annotated with the name of the talk and the speaker. Table 2 lists statistics on the three language pairs.

This data is made available under the Creative Commons license, Attribution-Non Commercial-No Derivatives (or the CC BY-NC-ND 4.0 International, <https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode>), all credit for the content goes to the TED organization and the respective authors of the talks. The data itself can be found at <http://www.cs.cmu.edu/~pmichell/sated/>.

## 5 Experiments

We run a set of experiments to validate the ability of our proposed approach to model speaker-induced variations in translation.

### 5.1 Experimental setup

We test three models base (a baseline ignoring speaker labels), full\_bias and fact\_bias. During training, we limit our vocabulary to the 40,000 most frequent words. Additionally, we discard any word appearing less than 2 times. Any word that doesn't satisfy those conditions is replaced with an UNK token.<sup>5</sup>

All our models are implemented with the DyNet (Neubig et al., 2017) framework, and unless specified we use the default settings therein. We refer to appendix B for a detailed explanation of the training process. We translate the test set using beam search with beam size 5.

### 5.2 Does explicitly modeling speaker-related variation improve translation quality?

Table 3 shows final test scores for each model with statistical significance measured with paired boot-

<sup>5</sup>Recent NMT systems also commonly use sub-word units (Sennrich et al., 2016b). This may influence on the result, either negatively (less direct control over high-frequency words) or positively (more capacity to adapt to high-frequency words). We leave a careful examination of these effects for future work.

	en-fr	en-es	en-de
base	38.05	39.89	26.46
spk_token	<b>38.85</b>	40.04	26.52
full_bias	<b>38.54</b>	<b>40.30</b>	<b>27.20</b>
fact_bias	<b>39.01</b>	39.88	<b>26.94</b>

Table 3: Test BLEU. Scores significantly ( $p < 0.05$ ) better than the baseline are written in bold

strap resampling (Koehn, 2004). As shown in the table, both proposed methods give significant improvements in BLEU score, with the biggest gains in English to French (+0.99) and smaller gains in German and Spanish (+0.74 and +0.40 respectively). Reducing the number of parameters with fact\_bias gives slightly better (en-fr) or worse (en-de) BLEU score, but in those cases the results are still significantly better than the baseline.

However, BLEU is not a perfect evaluation metric. In particular, we are interested in evaluating how much of the personal traits of each speaker our models capture. To gain more insight into this aspect of the MT results, we devise a simple experiment. For every language pair, we train a classifier (continuous bag-of-n-grams; details in Appendix C) to predict the author of each sentence on the target language part of the training set. We then evaluate the classifier on the ground truth and the outputs from our 3 models (base, full\_bias and fact\_bias).

The results are reported in Figure 2. As can be seen from the figure, it is easier to predict the author of a sentence from the output of speaker-specific models than from the baseline. This demonstrates that explicitly incorporating information about the author of a sentence allows for better transfer of personal traits during translations, although the difference from the ground truth demonstrates that this problem is still far from solved. Appendix D shows qualitative examples of our model improving over the baseline.

### 5.3 Further experiments on the Europarl corpus

One of the quirks of the TED talks is that the speaker annotation correlates with the topic of their talk to a high degree. Although the topics that a speaker talks about can be considered as a manifestation of speaker traits, we also perform a control experiment on a different dataset to verify that our model is indeed learning more than just topical

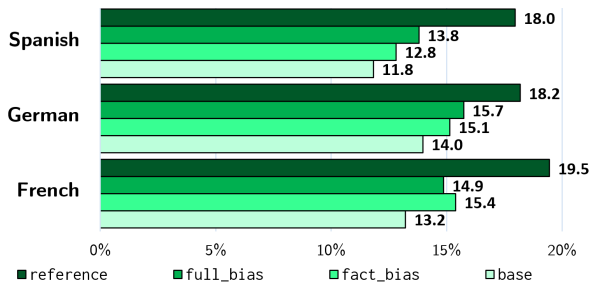


Figure 2: Speaker classification accuracy of our continuous bag-of-n-grams model.

information. Specifically, we train our models on a speaker annotated version of the Europarl corpus (Rabinovich et al., 2017), on the en-de language pair<sup>6</sup>.

We use roughly the same training procedure as the one described in §5.1, with a random train/dev/test split since none is provided in the original dataset. Note that in this case, the number of speakers is much lower (747) whereas the total size of the dataset is bigger ( $\approx 300k$ ).

We report the results in table 4. Although the difference is less salient than in the case of SATED, our factored bias model still performs significantly better than the baseline (+0.83 BLEU). This suggests that even outside the context of TED talks, our proposed method is capable of improvements over a speaker-agnostic model.

## 6 Related work

Domain adaptation techniques for MT often rely on data selection (Moore and Lewis, 2010; Li et al., 2010; Chen et al., 2017; Wang et al., 2017), tuning (Luong and Manning, 2015; Miceli Barone et al., 2017), or adding domain tags to NMT input (Chu et al., 2017). There are also methods that fine-tune parameters of the model on each sentence in the test set (Li et al., 2016), and methods that adapt based on human post-edits (Turchi et al., 2017), although these follow our baseline adaptation strategy of tuning all parameters. There are also partial update methods for transfer learning, albeit for the very different task of transfer between language pairs (Zoph et al., 2016).

Pioneering work by Mima et al. (1997) introduced ways to incorporate information about speaker role, rank, gender, and dialog domain for

<sup>6</sup>available here: <https://www.kaggle.com/ellarabi/europarl-annotated-for-speaker-gender-and-age/version/1>

	en-de
base	26.04
spk_token	26.49
full_bias	26.44
fact_bias	<b>26.87</b>

Table 4: Test BLEU on the Europarl corpus. Scores significantly ( $p < 0.05$ ) better than the baseline are written in bold

rule based MT systems. In the context of data-driven systems, previous work has treated specific traits such as politeness or gender as a “domain” in domain adaptation models and applied adaptation techniques such as adding a “politeness tag” to moderate politeness (Sennrich et al., 2016a), or doing data selection to create gender-specific corpora for training (Rabinovich et al., 2017). The aforementioned methods differ from ours in that they require explicit signal (gender, politeness...) for which labeling (manual or automatic) is needed, and also handle a limited number of “domains” ( $\approx 2$ ), where our method only requires annotation of the speaker, and must scale to a much larger number of “domains” ( $\approx 1,800$ ).

## 7 Conclusion

In this paper, we have explained and motivated the challenge of modeling the speaker explicitly in NMT systems, then proposed two models to do so in a parameter-efficient way. We cast this problem as an extreme form of domain adaptation and showed that, even when adapting a small proportion of parameters (the softmax bias,  $< 0.1\%$  of all parameters), allowed the model to better reflect personal linguistic variations through translation.

We further showed that the number of parameters specific to any person could be reduced to as low as 10 while still retaining better scores than a baseline for some language pairs, making it viable in a real world application with potentially millions of different users.

## Acknowledgements

The authors give their thanks the anonymous reviewers for their useful feedback which helped make this paper what it is, as well as the members of Neulab who helped proof read this paper and provided constructive criticism. This work was supported by a Google Faculty Research Award 2016 on Machine Translation.

## References

- Boxing Chen, Colin Cherry, George Foster, and Samuel Larkin. 2017. [Cost weighting for neural machine translation domain adaptation](#). In *Proceedings of the First Workshop on Neural Machine Translation*. Association for Computational Linguistics, Vancouver, pages 40–46. <http://www.aclweb.org/anthology/W17-3205>.
- Chenhui Chu, Raj Dabre, and Sadao Kurohashi. 2017. [An empirical comparison of domain adaptation methods for neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Vancouver, Canada, pages 385–391. <http://aclweb.org/anthology/P17-2061>.
- Michael Denkowski and Graham Neubig. 2017. Stronger baselines for trustable results in neural machine translation. *arXiv preprint arXiv:1706.09733*.
- Yarin Gal and Zoubin Ghahramani. 2016. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*. pages 1019–1027.
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. [Deep unordered composition rivals syntactic methods for text classification](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Beijing, China, pages 1681–1691. <http://www.aclweb.org/anthology/P15-1162>.
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. In *ICLR*.
- Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In Dekang Lin and Dekai Wu, editors, *Proceedings of EMNLP 2004*. Association for Computational Linguistics, Barcelona, Spain, pages 388–395.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the ACL on Interactive Poster and Demonstration Sessions*. Association for Computational Linguistics, Stroudsburg, PA, USA, ACL '07, pages 177–180. <http://dl.acm.org/citation.cfm?id=1557769.1557821>.
- Mu Li, Yingdong Zhao, Dongdong Zhang, and Ming Zhou. 2010. [Adaptive development data selection for log-linear model in statistical machine translation](#). In *Proceedings of the 23rd International Conference on Computational Linguistics (Coling 2010)*. Coling 2010 Organizing Committee, Beijing, China, pages 662–670. <http://www.aclweb.org/anthology/C10-1075>.
- Xiaoqing Li, Jiajun Zhang, and Chengqing Zong. 2016. One sentence one model for neural machine translation. *arXiv preprint arXiv:1609.06490*.
- Minh-Thang Luong and Christopher D Manning. 2015. Stanford neural machine translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*.
- Antonio Valerio Miceli Barone, Barry Haddow, Ulrich Germann, and Rico Sennrich. 2017. [Regularization techniques for fine-tuning in neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, pages 1490–1495. <https://www.aclweb.org/anthology/D17-1156>.
- Hideki Mima, Osamu Furuse, and Hitoshi Iida. 1997. Improving performance of transfer-driven machine translation with extra-linguistic information from context, situation and environment. In *IJCAI (2)*. pages 983–989.
- Shachar Mirkin, Scott Nowson, Caroline Brun, and Julien Perez. 2015. [Motivating personality-aware machine translation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, pages 1102–1108. <http://aclweb.org/anthology/D15-1130>.
- Robert C. Moore and William Lewis. 2010. [Intelligent selection of language model training data](#). In *Proceedings of the ACL 2010 Conference Short Papers*. Association for Computational Linguistics, Uppsala, Sweden, pages 220–224. <http://www.aclweb.org/anthology/P10-2041>.
- Graham Neubig, Chris Dyer, Yoav Goldberg, Austin Matthews, Waleed Ammar, Antonios Anastasopoulos, Miguel Ballesteros, David Chiang, Daniel Clothiaux, Trevor Cohn, et al. 2017. Dynet: The dynamic neural network toolkit. *arXiv preprint arXiv:1701.03980*.
- 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: Volume 2, Short Papers*. Association for Computational Linguistics, Valencia, Spain, pages 157–163. <http://www.aclweb.org/anthology/E17-2025>.
- Ella Rabinovich, Raj Nath Patel, Shachar Mirkin, Lucia Specia, and Shuly Wintner. 2017. [Personalized machine translation: Preserving original author traits](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume*

- 1, *Long Papers*. Association for Computational Linguistics, Valencia, Spain, pages 1074–1084. <http://www.aclweb.org/anthology/E17-1101>.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, San Diego, California, pages 35–40. <http://www.aclweb.org/anthology/N16-1005>.
- 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 (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 1715–1725. <http://www.aclweb.org/anthology/P16-1162>.
- Nitish Srivastava, Geoffrey E Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. 2014. Dropout: a simple way to prevent neural networks from overfitting. *Journal of machine learning research* 15(1):1929–1958.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 2818–2826.
- Marco Turchi, Matteo Negri, M Amin Farajian, and Marcello Federico. 2017. Continuous learning from human post-edits for neural machine translation. *The Prague Bulletin of Mathematical Linguistics* 108(1):233–244.
- Rui Wang, Masao Utiyama, Lemao Liu, Kehai Chen, and Eiichiro Sumita. 2017. Instance weighting for neural machine translation domain adaptation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, pages 1483–1489. <https://www.aclweb.org/anthology/D17-1155>.
- 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*. Association for Computational Linguistics, Austin, Texas, pages 1568–1575. <https://aclweb.org/anthology/D16-1163>.