

Multi-lingual neural title generation for e-Commerce browse pages

Prashant Mathur and Nicola Ueffing and Gregor Leusch

MT Science Team

eBay

Kasernenstraße 25

Aachen, Germany

Abstract

To provide better access of the inventory to buyers and better search engine optimization, e-Commerce websites are automatically generating millions of easily searchable browse pages. A browse page groups multiple items with shared characteristics together. It consists of a set of slot name/value pairs within a given category that are linked among each other and can be organized in a hierarchy. This structure allows users to navigate laterally between different browse pages (i.e. browse between related items) or to dive deeper and refine their search. These browse pages require a title describing the content of the page. Since the number of browse pages is huge, manual creation of these titles is infeasible. Previous statistical and neural generation approaches depend heavily on the availability of large amounts of data in a language. In this research, we apply sequence-to-sequence models to generate titles for high- & low-resourced languages by leveraging transfer learning. We train these models on multi-lingual data, thereby creating one joint model which can generate titles in various different languages. Performance of the title generation system is evaluated on three different languages; English, German, and French, with a particular focus on low-resourced French language.

1 Introduction

Natural language generation (NLG) has a broad range of applications, from question answering systems to story generation, summarization etc. In this paper, we target a particular use case that is important for e-Commerce websites, which group multiple items on common pages called *browse pages* (BP). Each browse page contains an overview of various items which share some characteristics expressed as slot/value pairs.

For example, we can have a browse page for Halloween decoration, which will display different types like lights, figurines, and candy bowls. These different items of decoration have their own browse pages, which are linked from the BP for Halloween decoration. A ceramic candy bowl for Halloween can appear on various browse pages, e.g. on the BP for Halloween decoration, BP for Halloween candy bowls, as well as the (non Halloween-specific) BP for ceramic candy bowls.

To show customers which items are grouped on a browse page, we need a human-readable title of the content of that particular page. Different combinations of characteristics bijectively correspond to different browse pages, and consequently to different browse page titles.

Note that here, different from other natural language generation tasks described in the literature, slot names are already given; the task is to generate a title for a set of slots. Moreover, we do not perform any selection of the slots that the title should realize; but all slots need to be realized in order to have a unique title. E-Commerce sites may have tens of millions of such browse pages in many different languages. The number of unique slot-value pairs are in the order of hundreds of thousands. All these factors render the task of human creation of BP titles infeasible.

Mathur, Ueffing, and Leusch (2017) developed several different systems which generated titles for these pages automatically. These systems include rule-based approaches, statistical models, and combinations of the two. In this work, we investigate the use of neural sequence-to-sequence models for browse page title generation. These models have recently received much attention in the research community, and are becoming the new state of the art in machine translation (refer Section 4).

We will compare our neural generation models

against two state-of-the-art systems.

1. The baseline system for English and French implements a *hybrid* generation approach, which combines a rule-based approach (with a manually created grammar) and statistical machine translation (SMT) techniques. For French, we have monolingual data for training language model, which can be used in the SMT system. For English, we also have human-curated titles and can use those for training additional “translation” components for this hybrid system.
2. The system for German is an *Automatic Post-Editing (APE)* system – first introduced by Simard et al. (2007) – which generates titles with the rule-based approach, and then uses statistical machine translation techniques for automatically correcting the errors made by the rule-based approach.

In the following section, we describe a few of the previous works in the field of language generation from a knowledge base or linked data. Section 3 addresses the idea of *lexicalization* of a browse node in linear form along with the *normalization* step to replace the slot values with placeholders. Sequence-to-sequence models for generation of titles are described in Section 4, followed by a description of joint learning over multiple languages in Section 5. Experiments and results are described in Sections 6 and 7.

2 Related work

The first works on NLG were mostly focused on rule-based language generation (Dale et al., 1998; Reiter et al., 2005; Green, 2006). NLG systems typically perform three different steps: *content selection*, where a subset of relevant slot/value pairs are selected, followed by *sentence planning*, where these selected pairs are realized into their respective linguistic variations, and finally *surface realization*, where these linguistic structures are combined to generate text. Our use case differs from the above in that there is no selection done on the slot/value pairs, but all of them undergo the sentence planning step. In rule-based systems, all of the above steps rely on hand-crafted rules.

Data driven approaches, on the other hand, either try to learn each of the steps automatically from the data Barzilay and Lapata (2005)

Dale et al. (1998) described the problem of generating natural language titles and short descriptions of structured nodes which consist of slot/value pairs. There are many research which deal with learning a generation model from parallel data. These parallel data consist of the structured data and natural-language text, so that the model can learn to transform the structured data into text. Duma and Klein (2013) generate short natural-language descriptions, taking structured DBpedia data as input. Their approach learns text templates which are filled with the information from the structured data.

Mei et al. (September, 2015) use recurrent neural network (LSTM) models to generate text from facts given in a knowledge base. Chisholm et al. (2017) solve the same problem by applying a machine translation system to a linearized version of the pairs. Several recent papers tackle the problem of generating a one-sentence introduction for a biography given structured biographical slot/value pairs. One difference between our work and the papers above, (Mei et al., September, 2015), and (Chisholm et al., 2017), is that they perform selective generation, i.e. they run a selection step that determines the slot/value pairs which will be included in the realization. In our use case however, all slot/value pairs are relevant and need to be realized.

Serban et al. (2016) generate questions from facts (structured input) by leveraging fact embeddings and then employing placeholders for handling rare words. In their work, the placeholders are heuristically mapped to the facts, however, we map our placeholders depending on the neural attention (for details, see Section 4).

3 Lexicalization

Our first step towards title generation is verbalization of all slot/value pairs. This can be achieved by a rule-based approach as described in (Mathur et al., 2017). However, in the work presented here, we do not directly lexicalize the slot/value pairs, but realize them in a pseudo language first. For example, the pseudo-language sequence for the slot/value pairs in Table 1 is “*_brand* ACME *_cat* Cell Phones & Smart Phones *_color* white *_capacity* 32GB”.¹

¹_cat refers to an e-Commerce category in the browse page.

Slot Name	Value
Category	<i>Cell Phones & Smart Phones</i>
Brand	<i>ACME</i>
Color	<i>white</i>
Storage Capacity	<i>32GB</i>

Table 1: Example of browse page slot/value pairs.

3.1 Normalization

Pseudo-language browse pages can still contain a large number of unique slot values. For example, there exist many different brands for smart phones (Samsung, Apple, Huawei, etc.). Large vocabulary is a known problem for neural systems, because rare or less frequent words tend to translate incorrectly due to data sparseness (Luong et al., 2015). At the same time, the softmax computation over the large vocabulary becomes intractable in current hardware. To avoid this issue, we normalize the pseudo-language sequences and thereby reduce the vocabulary size. For each language, we computed the 30 most frequent slot names and normalized their values via placeholders (Luong et al., August, 2015). For example, the lexicalization of “*Brand: ACME*” is “*_brand ACME*”, but after normalization, this becomes *_brand \$brand|ACME*. This representation means that the slot name *_brand* has the value of a placeholder *_brand* which contains the entity called “ACME”. During training, we remove the entity from the normalized sequence, while keeping them during translation of development or evaluation set. The mapping of placeholders in the target text back to entity names is described in Section 4.

The largest reduction in vocabulary size would be achieved by normalizing all slots. However, this would create several issues in generation. Consider the pseudo-language sequence “*_bike Road bike _type Racing*”. If we replace all slot values with placeholders, i.e. “*_bike \$bike _type \$type*”, then the system will not have enough information for generating the title “Road racing bike”. Moreover, the boolean slots, such as “*_comic Marvel comics _signed No*” would be normalized to placeholders as “*_comic \$comic _signed \$signed*”, and we would lose the information (“No”) necessary to realize this title as “Unsigned Marvel comics”.

3.2 Sub-word units

We applied another way of reducing the vocabulary, called byte pair encoding (BPE) (Sennrich

et al., 2016), a technique often used in NMT systems (Bojar et al., 2017). BPE is essentially a data compression technique which splits each word into sub-word units and allows the NMT system to train on a smaller vocabulary. One of the advantages of BPE is that it propagates generation of unseen words (even with different morphological variations). However, in our use case, this can create issues, because if BPE splits a brand and generates an incorrect brand name in the target, an e-Commerce company could be legally liable for the mistake. In such case, one can first run the normalization with placeholders followed by BPE, but due to time constraints, we do not report experiments on the same.

4 Sequence-to-Sequence Models

Sequence-to-sequence models in this work are based on an encoder-decoder model and an attention mechanism as described by Bahdanau et al. (May, 2016). In this network, the encoder is a bi-directional RNN which encodes the information of a sentence $X = (x_1, x_2, \dots, x_m)$ of length m into a fixed length vector of size $|h_i|$, where h_i is the hidden state produced by the encoder for token x_i . Since our encoder is a bi-directional model, the encoded hidden state is $h_i = h_{i, fwd} + h_{i, bwd}$, where h_{fwd} and h_{bwd} are unidirectional encoders, running from left to right and right to left, respectively. That is, they are encoding the context to the left and to the right of the current token.

Our decoder is a simple recurrent neural network (RNN) consisting of gated recurrent units (GRU) (Cho et al., 2014) because of their computationally efficiency. The RNN predicts the target sequence $Y = (y_1, y_2, \dots, y_j, \dots, y_l)$ based on the final encoded state h . Basically, the RNN predicts the target token $y_j \in V$ (with target vocabulary V) and emits a hidden state s_j based on the previous recurrent state s_{j-1} , the previous sequence of words $Y_{j-1} = (y_1, y_2, \dots, y_{j-1})$ and C_j , a weighted attention vector. The attention vector is a weighted average of all the hidden source states h_i , where $i = 1, \dots, m$. Attention weight (a_{ij}) is computed between the hidden states h_i and s_j and is leveraged as a weight of that source state h_i . In generation, we make use of these alignment scores to align our placeholders.² The target placeholders are bijectively mapped to those

²These placeholders are not to be confused with the placeholder for a tensor.

source placeholders whose alignment score (a_{ij}) is the highest at the time of generation.

The decoder predicts a score for all the tokens in the target vocabulary, which is then normalized by a softmax function, and the token with the highest probability is predicted.

5 Multilingual Generation

In this section, we present the extension of our work from a single-language setting to multi-language settings. There have been various studies in the past that target neural machine translation from multiple source languages into a single target language (Zoph and Knight, Jan, 2016), from single source to multiple target languages (Dong et al., 2015) and multiple source to multiple target languages (Johnson et al., June, 2016). One of the main motivation of joint learning in above works is to improve the translation quality on a low-resource language pair via transfer learning between related languages. For example, Johnson et al. (June, 2016) had no parallel data available to train a Japanese-to-Korean MT system, but training Japanese-English and English-Korean language pairs allowed their model to learn translations from Japanese to Korean without seeing any parallel data. In our case, the amount of training data for French is small compared to English and German (cf. Section 6.1). We propose joint learning of English, French and German, because we expect that transfer learning will improve generation for French. We investigate the joint training of pairs of these languages as well the combination of all three.

On top of the multi-lingual approach, we follow the work of Currey et al. (2017) who proposed copying monolingual data on both sides (source and target) as a way to improve the performance of NMT systems on low-resource languages. In machine translation, there are often named entities and nouns which need to be translated verbatim, and this copying mechanism helps in identifying them. Since our use case is monolingual generation, we expect a large gain from this copying approach because we have many brands and other slot values which need to occur verbatim in the generated titles.

6 Experiments

6.1 Data

We have access to a large number of human-created titles (curated titles) for English and German, and a small number of curated titles for French. When generating these titles, human annotators were specifically asked to realize all slots in the title.

We make use of a large monolingual out-of-domain corpus for French, as it is a low-resource language. We collect item description data from an e-Commerce website and clean the data in the following way: 1) we train a language model (LM) on the small amount of French curated titles, 2) we tokenize the out-of-domain data, 3) we remove all sentences with length less than 5, 4) we compute the LM perplexity for each sentence in the out-of-domain data, 5) we sort the sentences in increasing order of their perplexities and 6) select the top 500K sentences. Statistics of the data sets are reported in Table 2.

Languages	Set	#Titles	#trg Tokens
English	Train	222k	1.5M
	Dev	1000	6682
	Test	1000	6633
German	Train	226k	1.9M
	Dev	1000	8876
	Test	500	4414
French	Train	10k	95k
	Monolingual	500k	5.54M
	Dev	486	6403
	Test	478	3886

Table 2: Training and test data statistics per language. ‘k’ and ‘M’ stands for thousand and million, respectively.

6.2 Systems

We compared the NLG systems in the single-, dual-, and multi-lingual settings.

Single-language setting: This is the baseline NLG system, a straightforward sequence-to-sequence model with attention as described in Luong et al. (August, 2015), trained separately for each language. The vocabulary is computed on the concatenation of both source and target data, and the same vocabulary is used for both source and target languages in the experiments.

We use Adam (Kingma and Ba, December, 2014) as a gradient descent approach for faster convergence. Initial learning rate is set to 0.0002 with a decay rate of 0.9. The dimension of word embeddings is set to 620 and hidden layer size to

1000. Dropout is set to 0.2 and is activated for all layers except the initial word embedding layer, because we want to realize all aspects, we cannot afford to zero out any token in the source. We continue training of the model and evaluate on the development set after each epoch, stopping the training if the BLEU score on the development set does not increase for 10 iterations.

Baselines: We compare our neural system with a **fair** baseline system (*Baseline 1*), which is a statistical MT system trained on the same parallel data as the neural system: the source side is the linearized pseudo-language sequence, and the target side is the curated title in natural language. *Baseline 2* is either the hybrid system (for French and English) or the APE system (for German), both described in Section 1. These are **unfair** baselines, because (1) the hybrid system employs a large number of hand-made rules in combination with statistical models (Mathur, Ueffing, and Leusch, 2017), while the neural systems are unaware of the knowledge encoded in those rules, (2) the APE system and neural systems learn from same *amount* of parallel data, but the APE system aims at correcting rule-based generated titles, whereas the neural system aims at generating titles directly from a linearized form, which is a harder task. We compare our systems with the best performing systems of (Mathur et al., 2017), i.e. hybrid system for English and French, and APE system for German.

Multi-lingual setting: We train the neural model jointly on multiple languages to leverage transfer learning from a high-resource language to a low-resource one. In our multi-lingual setting, we experiment with three different combinations to improve models for French: 1) English+French (*en-fr*) 2) German+French (*de-fr*) 3) English+French+German (*en-fr-de*). English and French being close languages, we expect the *en-fr* system to benefit more from transfer learning across languages than any other combination. Although, as evident in Zoph and Knight (Jan, 2016), joint learning between the distant languages works better as they tend to disambiguate each other better than two languages which are close. For comparison, we also run a combination of two high-resource languages, i.e. English and German (*en-de*), to see if transfer learning works for them. It is important to note that in all multi-lingual sys-

tems the low-resourced language is over-sampled to balance the data.

We used the same design parameters on the neural network in both the single-language and the multi-lingual setting.

Normalized setting: On top of the systems above, we also experimented with the normalization scheme presented in Section 3.1. Normalization is useful in two ways: 1) It reduces the vocabulary size and 2) it avoids spurious generation of important aspect values (slot values). The second point is especially important in our case because this avoids highly sensitive issues such as brand violations. MT researches have observed that NMT systems often generate very fluent output, but have a tendency to generate inadequate output, i.e. sentences or words which are not related to the given input (Koehn and Knowles, June, 2017). We alleviate this problem through the normalization described above. After normalization, we see vocabulary reductions of 15% for French, 20% for German and as high as 35% for English.

As described in Section 5, we also use byte pair encoding, with a BPE code size of 30,000 for all systems (with BPE). We train the codes on the concatenation of source and target since (in this monolingual generation task) the vocabularies are very similar; the vocabulary size is around 30k for systems using BPE for both source and target.

7 Results

We evaluate our systems with three different automatic metrics: BLEU (Papineni et al., 2002), TER (Snover et al., 2006) and character F-Score (Popović, 2016). Note that BLEU and character F-score are quality metrics, i.e. higher scores mean higher quality, while TER is an error metric, where lower scores indicate higher quality. All metrics compare the automatically generated title against a human-curated title and determine sequence matches on the word or character level.

Table 3 summarizes results from all systems on the **English** test set. All neural systems are better than the fair *Baseline 1* system.

Normalization with tags (i.e. using placeholders) has a negative effect on English title quality both in the single-language setting *en* (67.1 vs. 68.4 BLEU) and in the dual-language setting *en-fr* (67.1 vs. 70.7 BLEU). However, title quality increases when using BPE instead (71.9 vs. 70.7 BLEU). On *en-de*, we observe gains

System	Norm.	BLEU↑	chrF1↑	TER↓
Baseline 1	n/a	64.2	82.9	26.5
Baseline 2	n/a	74.3	86.1	19.8
<i>en</i>	No	68.4	82.8	21.2
<i>en</i>	Yes(Tags)	67.1	82.5	21.7
<i>en-fr</i>	No	70.7	83.9	20.1
<i>en-fr</i>	Yes(Tags)	67.1	82.1	22.8
<i>en-fr</i>	Yes(BPE)	71.9	85.2	18.5
<i>en-fr_{big}</i>	Yes(BPE)	74.1	86.2	17.3
<i>en-de</i>	No	65.8	80.7	23.6
<i>en-de</i>	Yes(Tags)	67.1	82.8	22.3
<i>en-de</i>	Yes(BPE)	72.7	85.4	18.8
<i>en-fr-de</i>	Yes(BPE)	74.5	86.3	17.0

Table 3: Results on EN test, cased and detokenized.

both from normalization with tags and from BPE. Again, BPE normalization works best. Both dual-language systems with BPE achieve better performance than the best monolingual English system (71.9 and 72.7 vs. 68.4 BLEU).

The system *en-fr_{big}* contains monolingual French data added via the copying mechanism, which improves title quality. It outperforms any other neural system and is on par with *Baseline 2* (unfair baseline), even outperforming it in terms of TER. The multi-lingual system *en-fr-de* is very close to *en-fr_{big}* according to all three metrics.

System	Norm.	BLEU↑	chrF1↑	TER↓
Baseline 1	n/a	58.5	88.3	31.4
Baseline 2	n/a	79.4	90.7	17.1
<i>de</i>	No	78.2	87.0	20.7
<i>de</i>	Yes(Tags)	71.1	85.0	27.2
<i>en-de</i>	No	74.0	87.3	22.6
<i>en-de</i>	Yes(Tags)	65.6	84.0	30.2
<i>en-de</i>	Yes(BPE)	79.6	91.1	16.6
<i>de-fr</i>	No	77.2	88.9	18.9
<i>de-fr</i>	Yes(Tags)	63.3	83.0	30.7
<i>de-fr</i>	Yes(BPE)	77.6	89.0	19.2
<i>de-fr_{big}</i>	Yes(BPE)	80.0	91.6	16.2
<i>en-fr-de</i>	Yes(BPE)	80.6	92.0	15.3

Table 4: Results on DE test, cased and detokenized.

Table 4 collects the results for all systems on the **German** test set. For the single-language setting, we see a loss of 7 BLEU points when normalizing the input sequence, which is caused by incorrect morphology in the titles. When using placeholders, the system generates entities in the title in the exact form in which they occur in the input. In German, however, the words often need to be inflected. For example, the slot “*.brand Markenlos*” should be realized as “Markenlose” (Unbranded) in the title, but the placeholder generates the input form “Markenlos” (without suffix ‘e’). This causes a huge deterioration in the word-level met-

rics BLEU and TER, but not as drastic in chrF1, which evaluates on the character level.

For German, there is a positive effect of transfer learning for both dual-language systems *en-de* and *de-fr_{big}* with BPE (79.6 and 80.0 vs. 78.2 BLEU). However, the combination of languages hurts when we combine languages at token level, i.e. without normalization or with tags. The performance of systems with BPE is even on par with or better than the strong baseline of 79.4 BLEU, both for combinations of two and of three languages.

System	Norm.	BLEU↑	chrF1↑	TER↓
Baseline 1	n/a	44.6	77.7	44.3
Baseline 2	n/a	76.8	89.0	18.4
<i>fr_{small}</i>	No	23.0	52.0	71.1
<i>fr_{small}</i>	Yes(Tags)	27.4	56.2	60.1
<i>fr_{big}</i>	Yes(BPE)	29.5	57.3	58.5
<i>fr_{big}</i>	Yes(Both)	31.4	61.3	60.9
<i>en-fr</i>	No	22.5	51.3	69.6
<i>en-fr</i>	Yes(Tags)	20.1	47.1	70.3
<i>en-fr</i>	Yes(BPE)	21.6	50.7	73.9
<i>en-fr_{big}</i>	Yes(BPE)	32.6	61.8	51.2
<i>de-fr</i>	No	21.7	50.2	71.4
<i>de-fr</i>	Yes(Tags)	23.2	49.9	67.3
<i>de-fr</i>	Yes(BPE)	30.9	63.0	61.8
<i>de-fr_{big}</i>	Yes(BPE)	38.8	67.8	50.5
<i>en-fr-de</i>	Yes(BPE)	45.3	73.2	42.0

Table 5: Results on FR test, cased and detokenized.

System	Title
src	_cat Équipements de garage _brand Outifrance
ref	Équipements de garage Outifrance
<i>fr_{small}</i>	Équipements de suspension et de travail
<i>fr_{small,tags}</i>	Équipements de garage Outifrance
src	_cat Cylindres émetteurs d’embrayage pour automobiles _brand Vauxhall
ref	Cylindres émetteurs d’embrayage pour automobiles Vauxhall
<i>fr_{small}</i>	Perles d’embrayage pour automobile Vauxhall
<i>fr_{big}</i>	Cylindres émetteurs d’embrayage pour automobile Vauxhall
src	_cat Dessous de verre de table _brand Amadeus
ref	Dessous de verre de table Amadeus
<i>fr_{big}</i>	Guirlandes de verre Dunlop de table
<i>en-fr-de</i>	Dessous de verre de table Amadeus

Table 6: Examples from the french test set.

Table 5 summarizes the results from all systems on the **French** test set. The single-language *fr* NMT system achieves a low BLEU score compared to the SMT system Baseline 1 (23.0 vs. 44.6). This is due to the very small amount of parallel data, which is a setting where SMT typically outperforms NMT as evidenced in Zoph et al. (April, 2016). Normalization has a big positive impact on all French systems (e.g. 27.4

vs. 23.0 BLEU for *fr*).

The *de-fr* systems show a much larger gain from transfer learning than the *en-fr* systems, which validates Zoph and Knight (Jan, 2016)’s results, who show that transfer learning is better for distant languages than for similar languages.

For all three languages, copying monolingual data improves the NMT system by a large margin.

The multi-lingual *en-fr-de* (BPE) system (with copied monolingual data) is the best system for all three languages. It has the additional advantage of being one single model that can cater to all three languages at once.

Table 6 presents the example titles comparing different phenomena. The first block shows the usefulness of placeholders in system $fr_{small, tags}$ (i.e. fr_{small} , normalized with tags) where in comparison to fr_{small} the brand is generated verbatim. The second block shows the effectiveness of copying the data where “Cylindres” is generated correctly in the fr_{big} (with BPE) system in comparison to fr_{small} . The last block shows that reordering and adequacy in generation can be improved with the helpful signals from high-resourced English and German languages.

8 Conclusion

We developed neural language generation systems for an e-Commerce use case for three languages with very different amounts of training data and came to the following conclusions:

(1) The lack of resources in French leads to generation of low quality titles, but this can be drastically improved upon with transfer learning between French and English and/or German.

(2) In case of low-resource languages, copying monolingual data (even if out-of-domain) improves the performance of the system.

(3) Normalization with placeholders usually helps for languages with relatively easy morphology.

(4) It is important to over-sample the low-resourced languages in order to balance the high- & low-resourced data, thereby, creating a stable NLG system.

(5) For French, a low-resource language in our use case, the hybrid system which combines manual rules and SMT technology is still far better than the best neural system.

(6) The multi-lingual model has the best trade-off, as it achieves the best results among the neural

systems in all three languages and it is one single model which can be deployed easily on a single GPU machine.

Acknowledgments

Thanks to our colleague Pavel Petrushkov for all the help with the neural MT toolkit.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. May, 2016. Neural machine translation by jointly learning to align and translate. *CoRR* abs/1409.0473 [cs.CL].
- Regina Barzilay and Mirella Lapata. 2005. Collective content selection for concept-to-text generation. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Stroudsburg, PA, USA, HLT ’05, pages 331–338.
- Ondrej Bojar, Christian Buck, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, and Julia Kreutzer. 2017. Proceedings of the second conference on machine translation, volume 1: Research papers. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark.
- Andrew Chisholm, Will Radford, and Ben Hachey. 2017. Learning to generate one-sentence biographies from Wikidata. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. Association for Computational Linguistics, pages 633–642.
- Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. In *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Doha, Qatar, 25 October 2014*.
- Anna Currey, Antonio Valerio Miceli Barone, and Kenneth Heafield. 2017. Copied monolingual data improves low-resource neural machine translation. In *Proceedings of the Second Conference on Machine Translation*. Association for Computational Linguistics.
- Robert Dale, Stephen J Green, Maria Milosavljevic, Cécile Paris, Cornelia Verspoor, and Sandra Williams. 1998. The realities of generating natural language from databases. In *Proceedings of the 11th Australian Joint Conference on Artificial Intelligence*. pages 13–17.

- Daxiang Dong, Hua Wu, Wei He, Dianhai Yu, and Haifeng Wang. 2015. Multi-task learning for multiple language translation. In *ACL (1)*. The Association for Computer Linguistics.
- Daniel Duma and Ewan Klein. 2013. Generating natural language from linked data: Unsupervised template extraction. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Long Papers*. Association for Computational Linguistics, pages 83–94.
- Nancy Green. 2006. Generation of biomedical arguments for lay readers. In *Proceedings of the Fourth International Natural Language Generation Conference*. Association for Computational Linguistics, Stroudsburg, PA, USA, INLG '06.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhipeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. June, 2016. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *CoRR* abs/1611.04558 [cs.CL].
- Diederik P. Kingma and Jimmy Ba. December, 2014. Adam: A method for stochastic optimization. *CoRR* abs/1412.6980 [cs.LG].
- Philipp Koehn and Rebecca Knowles. June, 2017. Six challenges for neural machine translation. *CoRR* abs/1706.03872 [cs.CL].
- Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. August, 2015. Effective approaches to attention-based neural machine translation. *CoRR* abs/1508.04025 [cs.CL].
- Thang Luong, Ilya Sutskever, Quoc V. Le, Oriol Vinyals, and Wojciech Zaremba. 2015. Addressing the rare word problem in neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26-31, 2015, Beijing, China, Volume 1: Long Papers*. pages 11–19.
- Prashant Mathur, Nicola Ueffing, and Gregor Leusch. 2017. Generating titles for millions of browse pages on an e-commerce site. In *Proceedings of the International Conference on Natural Language Generation*.
- Hongyuan Mei, Mohit Bansal, and Matthew R. Walter. September, 2015. What to talk about and how? Selective generation using LSTMs with coarse-to-fine alignment. *Computing Research Repository (CoRR)* abs/1509.00838 [cs.CL].
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting on association for computational linguistics*. Association for Computational Linguistics, pages 311–318.
- Maja Popović. 2016. chrF deconstructed: beta parameters and n-gram weights. In *Proceedings of the First Conference on Machine Translation*. Association for Computational Linguistics, Berlin, Germany.
- Ehud Reiter, Somayajulu Sripada, Jim Hunter, and Ian Davy. 2005. Choosing words in computer-generated weather forecasts. *Artificial Intelligence* 167:137–169.
- 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 (Volume 1: Long Papers)*. Association for Computational Linguistics, pages 1715–1725.
- Iulian Vlad Serban, Alberto García-Durán, Çağlar Gülçehre, Sungjin Ahn, Sarath Chandar, Aaron C. Courville, and Yoshua Bengio. 2016. Generating factoid questions with recurrent neural networks: The 30m factoid question-answer corpus. 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*.
- Michel Simard, Cyril Goutte, and Pierre Isabelle. 2007. Statistical phrase-based post-editing. In *In Proceedings of NAACL*.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *In Proceedings of Association for Machine Translation in the Americas*. pages 223–231.
- Barret Zoph and Kevin Knight. Jan, 2016. Multi-source neural translation. *CoRR* abs/1601.00710 [cs.CL].
- Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. April, 2016. Transfer learning for low-resource neural machine translation. *arXiv:1604.02201 [cs.CL]*.