

Improving Retrieval Augmented Neural Machine Translation by Controlling Source and Fuzzy-Match Interactions

Cuong Hoang*, Devendra Sachan*, Prashant Mathur, Brian Thompson, Marcello Federico

AWS AI Labs

pramathu@amazon.com

Abstract

We explore zero-shot adaptation, where a general-domain model has access to customer or domain specific parallel data at inference time, but not during training. We build on the idea of Retrieval Augmented Translation (RAT) where top- k in-domain fuzzy matches are found for the source sentence, and target-language translations of those fuzzy-matched sentences are provided to the translation model at inference time. We propose a novel architecture to control interactions between a source sentence and the top- k fuzzy target-language matches, and compare it to architectures from prior work. We conduct experiments in two language pairs (En-De and En-Fr) by training models on WMT data and testing them with five and seven multi-domain datasets, respectively. Our approach consistently outperforms the alternative architectures, improving BLEU across language pair, domain, and number k of fuzzy matches with almost no trade-off on inference latency.

1 Introduction

Domain adaptation techniques such as fine-tuning (Freitag and Al-Onaizan, 2016; Luong and Manning, 2015) are highly effective at increasing in-domain performance of neural machine translation (NMT) systems, but are impractical in many realistic settings. For example, consider a single machine serving translations to thousands of customers, each with a private Translation Memory (TM). In this case, adapting, storing and loading large adapted models for each customer is computationally infeasible. In this paper we thus consider zero-shot adaptation instead, with a single general-domain model trained from heterogeneous sources that has access to the customer or domain specific TM only at inference time.

Our work builds on Retrieval Augmented Translation (RAT) (Li et al., 2022; Bulte and Tezcan, 2019; Xu et al., 2020; He et al., 2021; Cai et al., 2021), a paradigm which combines a translation model (Vaswani et al., 2017) with an external retriever module that retrieves the top- k most similar source sentences from a TM (i.e. "fuzzy matches") (Farajian et al., 2017; Gu et al., 2017; Bulte and Tezcan, 2019). The encoder encodes the input sentence along with the translations of the top- k fuzzy-matches and passes the resulting encodings to the decoder.

Prior RAT methods for NMT have fallen into two camps: Early work (Bulte and Tezcan, 2019; Zhang et al., 2018) concatenated the source sentence and the top- k fuzzy matches before encoding, relying on the encoder’s self-attention to compare the source sentence to each target sentences and determine which target phrases are relevant for the translation. More recent work (He et al., 2021; Cai et al., 2021) has opted to encode the source sentences and the top- k fuzzy matches *independently*, effectively shifting the entire burden of determining which target phrases are relevant to the decoder. We hypothesize that neither approach is ideal: In the first, the encoder has access to the information that we expect to be important (namely, the source and the fuzzy matches), but the self-attention also has potentially confusing/spurious connections. In the second, the encoder lacks the self-attention connections between the source and the fuzzy matches.

To address these issues, we propose a novel architecture which has self-attention connections between the source sentence and each fuzzy-match, but not between fuzzy-matches. We denote this method **RAT with Selective Interactions (RAT-SI)**. Our method is illustrated in Figure 1, along with two previously discussed approaches.

Experiments in five English-German (En-De) domain-specific test sets (Aharoni and Goldberg, 2020) and seven English-French (En-Fr) domain

*Work done while the authors were at AWS AI Labs.

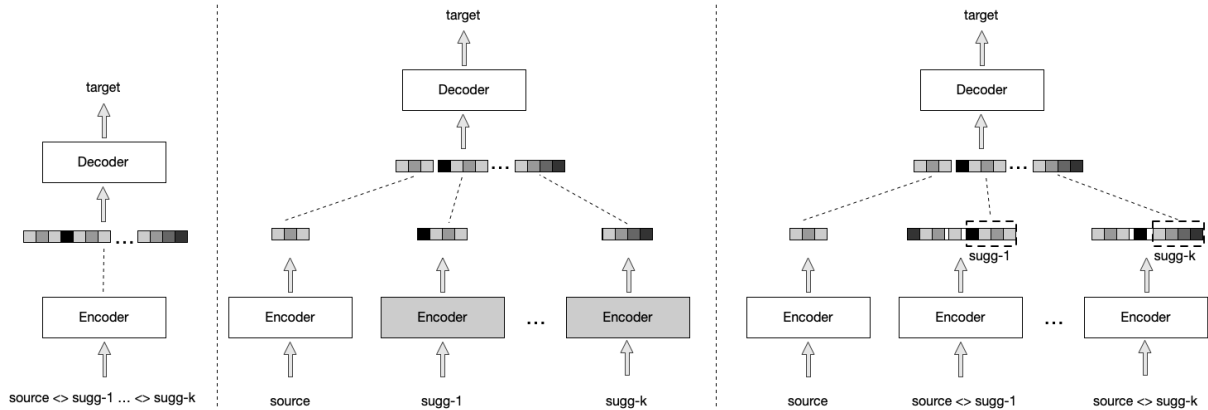


Figure 1: Architectures for retrieval augmented NMT. Left: Plain transformer ingesting source and retrieved fuzzy matches concatenated with a separator symbol (Bulte and Tezcan, 2019), denoted herein as RAT-CAT. Center: Transformer with dual encoder, one for encoding the source and one for encoding each retrieved fuzzy-match, inspired by He et al. (2021), denoted herein as RAT-SEP. Right: Transformer separately encoding the source and each source + fuzzy-match pair (this work), denoted herein as RAT-SI.

specific test sets (Pham et al., 2021), for $k = \{3, 4, 5\}$, demonstrate that our proposed method outperforms both prior approaches in 32 out of 36 cases considered. The proposed method outperforms the closest competitor by +0.82 to +1.75 BLEU for En-De and +1.57 to +1.93 for En-Fr.

2 Method

To isolate the effects of the underlying modeling strategy from the various tricks and implementation details employed in prior papers, we build baseline models which distill the two primary modeling strategies used in prior works:

The first concatenates a source sentence with target-language fuzzy matches and then encodes the entire sequence, as in Bulte and Tezcan (2019) and Xu et al. (2020). In this approach, the cross-attention of the encoder must learn to find the relevant parts of target-language fuzzy-matches by comparing each fuzzy-match to the source sentence, while ignoring potential spurious fuzzy-match to fuzzy-match interactions (see the left diagram in Figure 1). We denote this method **RAT-CAT**.

The second encodes the source and each target-language fuzzy-match separately (with two distinct encoders), and instead concatenates the encoded representations, inspired by He et al. (2021) and Cai et al. (2021). In this approach, the spurious connections between the target-language fuzzy-matches are eliminated, but the connections between the source and each fuzzy-match are also eliminated, forcing the attention in the decoder to

find the relevant portions in the fuzzy-match that are relevant to the source (see the center diagram in Figure 1). We denote this method **RAT-SEP**.

Finally, we propose a third method which attempts to build on the strengths of each of the prior methods. As in RAT-SEP, our method separately encodes (with the same encoder) the source and each target-language fuzzy-match; however, each fuzzy-match is jointly encoded with a copy of the source, as in RAT-CAT, allowing the encoder to find portions of the fuzzy-match which are relevant to the input. Finally, all the encoded inputs are concatenated and exposed to the decoder; However, the encoding of the source is only provided to the encoder once, to avoid potentially spurious interactions between copies of the input (see the right diagram in Figure 1). We denote our proposed method **RAT-SI**.

3 Experimental Setup

Our experiments are in two language directions: English-German (En-De) and English-French (En-Fr). We train models using the public WMT 2014 (Bojar et al., 2014) data set, with 4.5M En-De sentences and 36M En-Fr sentences.

During training, the model sees target-language fuzzy-match sentences from the same dataset it is being trained on (i.e. WMT14), but at inference, models must perform zero-shot adaptation to five En-De domain-specialized TMs¹ and seven En-Fr domain-specialized TMs.² En-De data is taken

¹Medical, Law, IT, Religion and Subtitles.

²News, Medical, Bank, Law, IT, TED and Religion.

from Aharoni and Goldberg (2020), which is a re-split version of the multi-domain data set from Koehn and Knowles (2017) while En-Fr data set is taken from the multi-domain data set of Pham et al. (2021).

To find target-language fuzzy matches for our model from domain specific TMs, we use Okapi BM25 (Robertson and Zaragoza, 2009), a classical retrieval algorithm that performs search by computing lexical matches of the query with all sentences in the evidence, to obtain top-ranked sentences for each input. To enable fast retrieval, we leverage the implementation provided by the Elasticsearch library.³ Specifically, we built an index using source sentences of each TM, and for every input source sentence, we collect top- k similar source side sentences and then use their corresponding target side sentences as inputs to the model.

To explore how each method performs (and how robust they are) under different conditions, we run a full set of experiments for $k = \{3, 4, 5\}$. We train separate models for each language pair and k value, and then apply that model to each of the 5 (En-De) or 7 (En-Fr) domains.

We report translation quality with BLEU scores computed via Sacrebleu (Post, 2018).⁴ We use compare-mt (Neubig et al., 2019) to perform pairwise significance testing with `bootstrap = 1000` and `prob_thresh = 0.05` for all pairs.

All models employed transformers (Vaswani et al., 2017) with 6 encoder and 6 decoder layers. Hidden size was set to 1024 and maximum input length truncated to 1024 tokens. All models employed a joint source-target language subword vocabulary of size $32K$ using Sentencepiece algorithm (Kudo and Richardson, 2018).

We use the Adam optimizer (Kingma and Ba, 2015) with $\beta_1 = 0.9$, $\beta_2 = 0.98$ and $\epsilon = 10^{-9}$; and (ii) increase the learning rate linearly for the first $4K$ training steps and decrease it thereafter; (iii) use batch size of $32K$ source tokens and $32K$ target tokens. Checkpoints are saved after every $10K$ iterations during training. We train models with maximum of $300K$ iterations. We use dropout of 0.1 and label-smoothing of 0.1.

³<https://github.com/elasticsearch/elasticsearch-py>

⁴SacreBleu signature: nrefs:1|case:mixed|eff:no|tok:13a|smooth:exp|version:2.0.0.

4 Results

Results for En-De are shown in Table 1 and results for En-Fr are shown in Table 2.

We observe several trends in the results. First, our proposed RAT-SI method outperforms both the RAT-CAT and RAT-SEP methods across both language pairs, having the best performance in 32/36 cases considered. In En-De, the proposed RAT-SI method has an average improvement of 1.43 BLEU over RAT-CAT and 2.35 BLEU over RAT-SEP, while in En-Fr we observe an average improvement of 1.73 BLEU over RAT-CAT and 2.98 over RAT-SEP. These results support our hypothesis that attention connections between the source sentences and each fuzzy match are critical to translation quality and the connections between the fuzzy matches are actually harmful.

Second, on average, $k = 5$ produces the best results for the RAT-SI method, but only by a small amount. However, considering individual language pair / domain combinations, there are many cases where $k = 5$ does not produce the best results, sometimes by several BLEU points. We hypothesize that this is due to the different domains containing, on average, different amounts of relevant data. This observation underscores the importance of tuning k , as well as testing new RAT methods under a variety of conditions, including different k values.

Finally, consistent with prior work, we see large improvements for all online domain-adapted methods (RAT-CAT, RAT-SEP, and RAT-SI) over the non-domain-adapted baseline, with improvements of up to +13.85 BLEU. This is not surprising, since the baseline model does not take advantage of any domain-specific data.

4.1 Latency

While not the focus of this work, we did a preliminary study of latency, comparing a baseline transformer to RAT-CAT and RAT-SI models during inference. We follow Domhan et al. (2020) and measure latency values as the 90th percentile of inference time when translating each sentence individually (no batching). We run experiments on an EC2 p3.2xlarge instance with a Tesla V100 GPU and report encoding latency results in Table 3. We use a batch size of 1 and $k=3$ for all experiments.

We observe a small increase of encoding latency by using RAT-SI (i.e. 17.48 ms) compared to of RAT-CAT. We provide a breakdown of the total

Model	k	IT	LAW	REL	MED	SUBT	Average
Baseline	n/a	27.92	35.59	11.26	30.74	19.46	24.99
RAT-CAT		33.97	50.34	25.14	45.05	19.89	34.88
RAT-SEP	k=3	32.78	49.04	22.92	44.28	20.48	33.90
RAT-SI (this work)		33.08	52.02*	26.40*	46.16	20.83*	35.70
RAT-CAT		33.67	49.59	23.40	44.87	20.27	34.36
RAT-SEP	k=4	31.84	48.38	24.37	43.55	19.99	33.63
RAT-SI (this work)		33.68	52.00*	28.42*	46.13	20.23	36.09
RAT-CAT		33.44	49.67	24.95	44.16	20.01	34.45
RAT-SEP	k=5	30.84	47.92	23.91	44.10	20.27	33.41
RAT-SI (this work)		33.84	52.17*	27.53*	46.95*	20.49	36.20

Table 1: BLEU scores for En-De experiments. The best BLEU for RAT models with a specific top-k value is **bolded**, and "*" indicates the best result is statistically significant compared to both the other methods. The proposed method (RAT-SI) produces the best results in 13/15 cases considered, with an average improvement of 1.43 BLEU over RAT-CAT and 2.35 BLEU over RAT-SEP.

Model	k	LAW	MED	IT	NEWS	BANK	REL	TED	Average
Baseline	n/a	52.68	31.12	32.22	35.09	41.04	14.51	35.55	34.60
RAT-CAT		66.32	37.09	39.91	35.09	49.01	61.83	36.39	46.52
RAT-SEP	k=3	64.93	37.06	38.02	35.57	49.14	53.34	37.29	45.05
RAT-SI (this work)		66.56	41.30*	40.61	35.53	50.19*	67.55*	37.42	48.45
RAT-CAT		65.71	37.31	38.71	34.60	49.43	63.55	36.10	46.49
RAT-SEP	k=4	64.35	37.89	38.89	35.45	49.28	53.41	37.13	45.20
RAT-SI (this work)		66.63*	39.50*	41.90*	35.71	50.04	65.20*	37.47	48.06
RAT-CAT		65.60	37.35	38.74	34.46	49.33	63.21	36.08	46.40
RAT-SEP	k=5	64.62	38.50	39.53	35.59	49.97	52.56	37.09	45.41
RAT-SI (this work)		67.03*	39.05	41.33*	35.93*	49.82	65.49*	37.90*	48.08

Table 2: BLEU scores for En-Fr experiments. The best BLEU for RAT models with a specific top-k value is **bolded**, and "*" indicates the best result is statistically significant compared to both the other methods. The proposed method (RAT-SI) produces the best results in 19/21 cases considered, with average improvements of 1.73 BLEU over RAT-CAT and 2.98 over RAT-SEP.

Model	Encoding Latency
Transformer	14.80 ms
RAT-CAT	15.23 ms
RAT-SI	17.48 ms

Table 3: Encoding latency in milliseconds of models (lower is better).

encoding latency in Table 4 which shows encoding the inputs in RAT-SI is faster than RAT-CAT but it requires an extra overhead for extracting the encoding of fuzzy matches from the joint encoding of source with fuzzy match. However, the encoding time is a very small fraction of overall latency (see Table 5) and thus this difference appears to be negligible.

We find that RAT-CAT and RAT-SI have nearly identical latencies, and each is only slightly slower than the baseline transformer (see Table 5). This is somewhat surprising since both methods make the input to the decoder significantly longer. We hypothesize that we are under-utilizing the GPU in all cases, and thus the increased computations does not increase latency. Further investigation of this is

RAT-SI Model	Encoding Latency
Encoding input	14.91 ms
Extra overheads	2.57 ms
Total time	17.48 ms

Table 4: Encoding latency of RAT-SI in milliseconds (lower is better). Extra overheads include (1): Concatenate input and $k = 3$ input-suggestion pairs (2): Extract $k = 3$ suggestion encodings and append them to the input encoding.

left for future work.

5 Related Work

Bulte and Tezcan (2019) proposed augmenting the input to NMT with target-language fuzzy-match sentences from a TM, concatenating the input and fuzzy-matches together. Their method was simpler than prior works such as (Zhang et al., 2018), which manipulated n-gram probabilities based on their occurrence in the fuzzy-matches. Xu et al. (2020) proposed several enhancements using the same architecture, including fine-tuning models

Model	Translation Latency
Transformer	574.02 ms
RAT-CAT	597.28 ms
RAT-SI	597.41 ms

Table 5: Translation latency in milliseconds of RAT-CAT and our model RAT-SI (lower is better). Batch size was set to one to simulate an on-demand system.

and masking out or marking words not related to the input sentence, and matching arbitrarily large n -grams instead of sentences.

More recent work has explored using separate encoders for input and fuzzy-match (He et al., 2021; Cai et al., 2021). He et al. (2021) also considers the realistic scenario where a TM may include noise, while Cai et al. (2021) explores finding target sentences in monolingual data instead of relying on a TM at inference time.

Xia et al. (2019) and Xu et al. (2020) explore aspects of filtering fuzzy-matches by applying similarity thresholds, leveraging word alignment information (Zhang et al., 2018; Xu et al., 2020; He et al., 2021) or re-ranking with additional score (e.g. word overlapping) (Gu et al., 2018; Zhang et al., 2018).

Our work is related to the use of k -nearest-neighbor for NMT (Khandelwal et al., 2021; Zheng et al., 2021) but it is less expensive and does not require storage and search over a large data store of context representations and corresponding target tokens (Meng et al., 2021).

Other works have considered online adaptation outside the context of RAT, including Vilar (2018), who proposes Learning Hidden Unit Contributions (Swietojanski et al., 2016) as a compact way to store many adaptations of the same general-domain model. For an overview of fuzzy-match augmentation outside of NMT, see Li et al. (2022).

Domain adaptation can also be performed offline, typically via fine tuning (Luong and Manning, 2015). Regularization is often applied during fine tuning to avoid catastrophic forgetting (Khayrallah et al., 2018; Thompson et al., 2019a,b).

TMs are commonly used in the localization industry to provide suggestions to translators in order to boost their productivity (Federico et al., 2012). Enhancing translation quality of MT system by leveraging fuzzy-matches extracted from TMs has been explored widely for statistical MT (Koehn and Senellart, 2010; Mathur et al., 2013) and neural MT

systems (Farajian et al., 2017; Gu et al., 2017; Cao and Xiong, 2018; Bulte and Tezcan, 2019).

6 Conclusion

Previous work in retrieval augmented translation has used architectures which either have full connections between source and all fuzzy matches, or independently encode the source and each fuzzy match. Based on our hypothesis that the attention connections between source and each fuzzy match are helpful, but that the connections between different fuzzy matches are harmful, we propose a new architecture (RAT-SI) with the former connections but not the latter. Experiments on several language pairs, domains, and different numbers of fuzzy matches (k) demonstrate that RAT-SI substantially outperforms the prior architectures.

7 Limitations

Due to the availability of domain specific datasets, we perform experiments on two high-resource languages, both out of English. It is unclear if our conclusions would hold on low-resource language pairs. Furthermore, our domains may or may not match real world use cases where an MT customer has their own TM. Real TMs may be significantly larger/smaller, contain multiple domains, etc.

References

- Roe Aharoni and Yoav Goldberg. 2020. [Unsupervised domain clusters in pretrained language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.
- Ondrej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, Radu Soricut, Lucia Specia, and Aleš Tamchyna. 2014. [Findings of the 2014 workshop on statistical machine translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58, Baltimore, Maryland, USA. Association for Computational Linguistics.
- Bram Bulte and Arda Tezcan. 2019. [Neural fuzzy repair: Integrating fuzzy matches into neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1800–1809, Florence, Italy. Association for Computational Linguistics.
- Deng Cai, Yan Wang, Huayang Li, Wai Lam, and Lema Liu. 2021. [Neural machine translation with monolingual translation memory](#). In *Proceedings*

- of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 7307–7318, Online. Association for Computational Linguistics.
- Qian Cao and Deyi Xiong. 2018. [Encoding gated translation memory into neural machine translation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3042–3047, Brussels, Belgium. Association for Computational Linguistics.
- Tobias Domhan, Michael Denkowski, David Vilar, Xing Niu, Felix Hieber, and Kenneth Heafield. 2020. [The sockeye 2 neural machine translation toolkit at AMTA 2020](#). In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 110–115, Virtual. Association for Machine Translation in the Americas.
- M. Amin Farajian, Marco Turchi, Matteo Negri, and Marcello Federico. 2017. [Multi-domain neural machine translation through unsupervised adaptation](#). In *Proceedings of the Second Conference on Machine Translation*, pages 127–137, Copenhagen, Denmark. Association for Computational Linguistics.
- Marcello Federico, Alessandro Cattelan, and Marco Trombetti. 2012. [Measuring user productivity in machine translation enhanced computer assisted translation](#). In *Proceedings of the 10th Conference of the Association for Machine Translation in the Americas: Research Papers*, San Diego, California, USA. Association for Machine Translation in the Americas.
- Markus Freitag and Yaser Al-Onaizan. 2016. Fast domain adaptation for neural machine translation. *arXiv preprint arXiv:1612.06897*.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O. K. Li. 2017. [Search engine guided non-parametric neural machine translation](#). *CoRR*, abs/1705.07267.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O. K. Li. 2018. [Search engine guided neural machine translation](#). In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, pages 5133–5140. AAAI Press.
- Qiuxiang He, Guoping Huang, Qu Cui, Li Li, and Lemao Liu. 2021. [Fast and accurate neural machine translation with translation memory](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3170–3180, Online. Association for Computational Linguistics.
- Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2021. [Nearest neighbor machine translation](#). In *International Conference on Learning Representations*.
- Huda Khayrallah, Brian Thompson, Kevin Duh, and Philipp Koehn. 2018. [Regularized training objective for continued training for domain adaptation in neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 36–44, Melbourne, Australia. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- 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, Vancouver. Association for Computational Linguistics.
- Philipp Koehn and Jean Senellart. 2010. [Convergence of translation memory and statistical machine translation](#). In *Proceedings of the Second Joint EM+/CNGL Workshop: Bringing MT to the User: Research on Integrating MT in the Translation Industry*, pages 21–32, Denver, Colorado, USA. Association for Machine Translation in the Americas.
- Taku Kudo and John Richardson. 2018. [Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 66–71. Association for Computational Linguistics.
- Huayang Li, Yixuan Su, Deng Cai, Yan Wang, and Lemao Liu. 2022. [A survey on retrieval-augmented text generation](#).
- Minh-Thang Luong and Christopher Manning. 2015. [Stanford neural machine translation systems for spoken language domains](#). In *Proceedings of the 12th International Workshop on Spoken Language Translation: Evaluation Campaign*, pages 76–79, Da Nang, Vietnam.
- Prashant Mathur, Mauro Cettolo, and Marcello Federico. 2013. [Online learning approaches in computer assisted translation](#). In *Proceedings of the Eighth Workshop on Statistical Machine Translation*, pages 301–308, Sofia, Bulgaria. Association for Computational Linguistics.
- Yuxian Meng, Xiaoya Li, Xiayu Zheng, Fei Wu, Xiaofei Sun, Tianwei Zhang, and Jiwei Li. 2021. [Fast nearest neighbor machine translation](#). *CoRR*, abs/2105.14528.

- Graham Neubig, Zi-Yi Dou, Junjie Hu, Paul Michel, Danish Pruthi, and Xinyi Wang. 2019. [compare-mt: A tool for holistic comparison of language generation systems](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 35–41, Minneapolis, Minnesota. Association for Computational Linguistics.
- MinhQuang Pham, Josep Maria Crego, and François Yvon. 2021. [Revisiting Multi-Domain Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 9:17–35.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Stephen Robertson and Hugo Zaragoza. 2009. [The probabilistic relevance framework: Bm25 and beyond](#). *Found. Trends Inf. Retr.*, 3(4):333–389.
- Pawel Swietojanski, Jinyu Li, and Steve Renals. 2016. Learning hidden unit contributions for unsupervised acoustic model adaptation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 24(8):1450–1463.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019a. [Overcoming catastrophic forgetting during domain adaptation of neural machine translation](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2062–2068, Minneapolis, Minnesota. Association for Computational Linguistics.
- Brian Thompson, Rebecca Knowles, Xuan Zhang, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019b. [HABLex: Human annotated bilingual lexicons for experiments in machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1382–1387, Hong Kong, China. 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 30*, pages 5998–6008. Curran Associates, Inc.
- David Vilar. 2018. [Learning hidden unit contribution for adapting neural machine translation models](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 500–505, New Orleans, Louisiana. Association for Computational Linguistics.
- Mengzhou Xia, Guoping Huang, Lemao Liu, and Shuming Shi. 2019. [Graph based translation memory for neural machine translation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7297–7304.
- Jitao Xu, Josep Crego, and Jean Senellart. 2020. [Boosting neural machine translation with similar translations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1580–1590, Online. Association for Computational Linguistics.
- Jingyi Zhang, Masao Utiyama, Eiichiro Sumita, Graham Neubig, and Satoshi Nakamura. 2018. [Guiding neural machine translation with retrieved translation pieces](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1325–1335, New Orleans, Louisiana. Association for Computational Linguistics.
- Xin Zheng, Zhirui Zhang, Junliang Guo, Shujian Huang, Boxing Chen, Weihua Luo, and Jiajun Chen. 2021. [Adaptive nearest neighbor machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 368–374, Online. Association for Computational Linguistics.