

Interrogating the Explanatory Power of Attention in Neural Machine Translation

Pooya Moradi, Nishant Kambhatla, and Anoop Sarkar

Simon Fraser University

8888 University Drive

Burnaby, BC, Canada

{pooya_moradi, nkambhat, anoop}@sfu.ca

Abstract

Attention models have become a crucial component in neural machine translation (NMT). They are often implicitly or explicitly used to justify the model’s decision in generating a specific token but it has not yet been rigorously established to what extent attention is a reliable source of information in NMT. To evaluate the explanatory power of attention for NMT, we examine the possibility of yielding the same prediction but with counterfactual attention models that modify crucial aspects of the trained attention model. Using these counterfactual attention mechanisms we assess the extent to which they still preserve the generation of function and content words in the translation process. Compared to a state of the art attention model, our counterfactual attention models produce 68% of function words and 21% of content words in our German-English dataset. Our experiments demonstrate that attention models by themselves cannot reliably explain the decisions made by a NMT model.

¹

1 Introduction

One shortcoming of neural machine translation (NMT), and neural models in general, is that it is often difficult for humans to comprehend the reasons why the model is making predictions (Feng et al., 2018; Ghorbani et al., 2019). The main cause of such a difficulty is that in neural models, information is implicitly represented by real-valued vectors, and conceptual interpretation of these vectors remains a challenge. Why do we want neural models to be interpretable? In order to debug a neural model during the error analysis process in research experiments, it is necessary to know how much each part of the model is

¹The source code to reproduce the experiments is available at: https://github.com/sfu-natlang/attention_explanation

contributing to the error in the prediction. Being able to interpret the deficiencies of a model is also crucial to further improve upon it. This requires an explainable understanding of the internals of the model, including how certain concepts are being modeled or represented. Therefore developing methods to interpret and understand neural models is an important research goal.

Visualizing and interpreting neural models has been extensively studied in computer vision (Simonyan et al., 2013; Bach et al., 2015; Zeiler and Fergus, 2014; Montavon et al., 2017), and more recently in natural language processing (NLP) (Karpathy et al., 2015; Li et al., 2016a; Strobel et al., 2017, 2018). Recently, the integration of attention mechanism (Bahdanau et al., 2014) with an NMT sequence to sequence model (Sutskever et al., 2014) has led to significant improvements in translation quality especially for longer sentence lengths. The attention mechanism provides a weighted average over the information from the source encodings to be used at each translation step. These weights are often regarded as a measure of importance, and are implicitly or explicitly used as an explanation for the model’s decision. However it has not yet been established to what extent such explanations are reliable.

The example in Figure 1 shows a model translating the German sentence “und wir wollen dieses material für alle verfügbar machen.” (*and we want to make this material accessible to everyone.*). At the time that the model is translating “verfügbar” to “accessible”, it is mostly attending to “verfügbar” (left heatmap). It is tempting to conclude that “verfügbar” having the most attention is why the model is generating the token “accessible”. However, we manipulate the attention weights such that the “verfügbar” receives no attention and “alle” is given the most attention (right heatmap). We observe that in the second case, the




Figure 1: Two distinct attention weights yielding the same prediction. The model is translating the source word “verfügbar” to “accessible”. In the left attention heatmap, the focus of the attention is on the word “verfügbar”. However, in the right heatmap, “verfügbar” is not attended to at all and “alle” has received the most attention.

model makes the same decision again and outputs “accessible”. This example shows that using attention weights to reason about the model’s predictions can be misleading, as these two heatmaps convey different explanations.

There are relatively few previous studies on investigating the power of attention mechanism in rationalizing a model’s predictions in NLP (Jain and Wallace, 2019; Serrano and Smith, 2019) and they all target text classification tasks where attention is over the input document. Their findings do not generalize easily to NMT due to the difference in how the decoder works in the translation task which produces a sequence rather than a class label. NMT is a sequence-to-sequence (Seq2Seq) task, which is different from text classification. Also, the size of the output space is quite limited in text classification, whereas in NMT, it is equal to the vocabulary size of the target language that can be very large. Furthermore, different neural architectures (e.g., presence of an encoder and decoder in NMT), require different analysis and interpretations. To the best of our knowledge there is no existing work on addressing the question of interpretability of attention models for the machine translation task.

To investigate if the explanation implied by attention weights faithfully justifies why a model produced its output, we study to what extent it is possible to yield the same prediction by using counterfactual attention weights that suggest a contradictory explanation. The intuition behind this approach is that multiple contradictory justifications should not exist for a decision. To be specific, in the setting of an encoder-decoder model (Sec 2), we propose several counterfactual attention weighting methods that suggest different explanations compared to the original attention weights (Sec 3) and analyze their performance in

preserving the generation of function² and content words (Sec 4). Function words (e.g., *a*, *the*, *is*) have little lexical meaning in contrast to content words and thus we are curious whether explanatory power of the attention mechanism differs for generation of these two groups of words.

2 Encoder-Decoder Model with Attention Mechanism

Given a training sentence pair (\mathbf{x}, \mathbf{y}) where $\mathbf{x} = [x_1, x_2, \dots, x_m]$ is a sentence in the source language and $\mathbf{y} = [y_1, y_2, \dots, y_n]$ is its corresponding translation in the target language, the encoder, which is a recurrent neural network (RNN), runs over the source sentence to calculate the contextualized representation of the source words. Here, we use a bidirectional encoder, and concatenate the forward and backward hidden states to build the final representation.

$$\begin{aligned} \vec{h}_t &= f_{enc}(x_t, \vec{h}_{t-1}) \\ \overleftarrow{h}_t &= \overleftarrow{f}_{enc}(x_t, \overleftarrow{h}_{t+1}) \\ h_t &= [\vec{h}_t, \overleftarrow{h}_t] \end{aligned} \quad (1)$$

Then the decoder starts to generate output tokens using the following probability distribution:

$$p(y_t | y_{<t}, x) = \text{softmax}(g_{dec}(s_t, c_t))$$

with g_{dec} being a transformation function that produces a vocabulary-sized vector, and s_t is the hidden unit of the decoder’s RNN updated as:

$$s_t = f_{dec}(y_{t-1}, s_{t-1}, c_{t-1})$$

²The reference for function words (we added new function words including the EOS token to this) can be found at: semanticsimilarity.files.wordpress.com/2013/08/jim-oshea-fwlist-277.pdf

where f_{dec} is a RNN. Here c_t is the context vector calculated by attention mechanisms:

$$c_t = \sum_{i=1}^m \alpha_{ti} h_i$$

where α_{ti} is the normalized attention weights over the source context:

$$\alpha_{ti} = \frac{e^{a(s_t, h_i)}}{\sum_j e^{a(s_t, h_j)}}$$

Here, a is a scoring function that determines the contribution of each source context vector to the final context vector. Implementation of a depends on the choice of the attention method. In this work, we use general attention (Luong et al., 2015) as the scoring function:

$$a(s_t, h_i) = s_t^\top W_a h_i$$

3 Approach

Given a trained NMT model M , a test sentence \mathbf{x} with \mathbf{y} being its translation generated by M at the decoding step t in which α_t is the attention vector, attending to the source word at position $m_t = \operatorname{argmax}_i \alpha_t[i]$ or k -best attended-to words in the source are often implicitly or explicitly regarded as a justification for the model’s prediction at the time step t .

A vital criteria for such a justification is that one should not be able to find a contradictory explanation for the model’s decision. More precisely, if at the time of inference, it is possible to manipulate the original attention weights to constitute an alternative attention vector α'_t , such that $\operatorname{argmax}_i \alpha'_t[i] \neq m_t$ and the decision of the model is preserved, then these weights cannot be used for justification as they are contradictory. Thus, we are interested in assessing for what percentage of the words in the translation, counterfactual attention weights exist. These percentages can shed light on the reliability of the attention weights as a potential explanation. Note that at the inference time we manipulate attention vector for each output token separately and in isolation to make sure the output tokens at time steps $t + 1$ and after will not be affected by the change at time step t . This means that our output translations are unaffected by our counterfactual attention models which are purely for the examination of how attention might explain the model’s decision.

The main task here is to find a counterfactual attention vector if it exists at all. An exhaustive search approach in which every possible attention vector is examined is computationally intractable and unlikely to provide much insight. Instead we experiment with a few specific counterfactual attention weighting methods that we think provide the most insight. It is important to note that in this case, the calculated percentage of the preserved words will be a lower-bound estimation for the true percentage. As a result, the explanation offered by attention weights are more unreliable than what we’ll find. We experiment with the following attention methods to create counterfactual attention weights:

- **RandomPermute (Jain and Wallace, 2019):** We set $\alpha'_t = \text{random_permute}(\alpha_t)$ such that $\operatorname{argmax}_i \alpha'_t[i] \neq m_t$.
- **Uniform:** In this method, $\alpha'_t = \frac{1}{m} \vec{1}$. Here, m is the length of the source sentence.
- **ZeroOutMax:** A simple approach to create a counterfactual attention vector is to remove the maximum attention weight. So we set $a(s_t, h_{m_t}) = -\infty$.

We also experiment with four additional attention methods where our motivation is not finding counterfactual attention weights, but to improve our understanding of how attention weights influence the model’s prediction:

- **ZeroOut:** In what conditions does the decoder overlook the information provided by attention mechanism? To answer this question, we set all attention weights to zero at inference time (attention is still used while training the model).
- **LastEncoderState:** Here, we only use the final hidden state of the encoder as the context vector to be used in the decoder. Note that this is different from seq2seq without attention in which final hidden state of the encoder is used to initialize the decoder. When the focus is on the final hidden state of the encoder in the original attention weights, this method does not produce a counterfactual attention vector, which is why we don’t intend to use this method to create a contradictory explanation, but rather to gain more insight into the sensitivity of the model to attention weights.

- **OnlyMax**: In this and the following method, the source hidden state with the maximum attention still receives the highest attention, and so these two methods do not output counterfactual attention vectors. However we are curious to what extent other attention weights are required to preserve the model’s output. Note that the weights produced by these methods can be counted as contradictory when multiple attention weights are used for justifying predictions. Because although the most attended source context is not changed, the relative ranking of the rest of the source context in terms of attention weights is changed. However, these kinds of justifications are mostly discussed in text classification. In this specific method we only keep the most attended source hidden state: $\alpha'_t[m_t]=1$.
- **KeepMaxUniformOthers**: Here, we set $\alpha'_t[m_t] = \alpha_t[m_t]$, but for all other positions $\alpha'_t[i] = (1 - \alpha'_t[m_t])/m$. This is to investigate if using other source hidden states uniformly has any added benefit.

4 Experiments

4.1 Data

In this work we use the German-English dataset from IWSLT2014³. We concatenate dev2010, dev2012, tst2010, tst2011 and tst2012 to be used as the test data. Data is tokenized using Moses (Koehn et al., 2007).

4.2 Model Details

OpenNMT (Klein et al., 2017) is used for our Seq2Seq implementation. We use Long Short-Term Memory (LSTM) as RNN units. Each LSTM unit has 2 layers, and the dimension size for LSTM units and word embeddings is set to 500. The model is trained using Adam trainer with learning rate 0.001 for 50000 steps using early stopping. Vocabulary size for both the source and target language is set to 50000. Sentences longer than 50 tokens are pruned.

5 Results

Table 1 shows the percentage of function and content words generated by the trained model. As expected, the majority of the generated tokens are

³<https://sites.google.com/site/iwsltevaluation2014/>

function words. We discuss our findings in more detail in the subsequent subsections.

Number of tokens (+EOS)	139465
Percentage of function words	68%
Percentage of content words	32%

Table 1: Percentage of function and content words in the generated translation.

	<i>Method</i>	<i>% for FWs</i>	<i>% for CWs</i>
1	RandomPermute	33%	6%
2	Uniform	53%	11%
3	ZeroOutMax	52%	15%
4	Aggregate(1+2+3)	68%	21%
5	ZeroOut	9%	0%
6	LastEncoderState	20%	2%
7	OnlyMax	71%	83%
8	KeepMaxUniformOthers	86%	86%

Table 2: Percentage of the preserved function and content words in the proposed attention methods: Trying out all the methods to find a counterfactual attention vector maximizes the chance of success. We use methods in row 5-8 only to shed light on the sensitivity of the model’s output to perturbation in attention weight. They are not necessarily counted as counterfactual attention methods. Higher preservation rate stands for better performance.

5.1 Effectiveness of the proposed counterfactual attention methods

Table 2 shows the percentage of function and content words for which counterfactual attention weights were found using the proposed attention methods. The `Uniform` method (row 2) is the most effective method to create counterfactual attention weights for function words. However, for content words, the `ZeroOutMax` method (row 3) is the most successful method.

From Table 2, we also derive that `RandomPermute` is not as effective as the `Uniform` and `ZeroOutMax` methods. Our justification is that in the `RandomPermute` method, it is highly probable that the context vector is biased toward a random source hidden state. Such bias can lead to misleading noise in the context vector. However, there isn’t such a bias in the `Uniform` or `ZeroOutMax` methods.

To maximize the chance of finding a counterfactual attention, for each output token, we try out all the proposed methods to check if we can find a counterfactual attention (row 4). As evident from Table 2, this approach greatly increases the chance of finding a counterfactual attention. Note that as previously stated, these percentages are a lower-bound for the true percentage.

5.2 Function words are more easily generated compared to content words

An important observation in Table 2 is that the proposed methods are considerably more effective in preserving function words compared to content words. The production of function words rely more on the target context, in contrast to content words which rely more on the source context. Accordingly, perturbation in the original attention weights likely has significantly more impact on diminishing content words compared to function words.

This ties well with the main idea behind context gates in which the influence of source context and target context is controlled dynamically (Tu et al., 2017). Since the generation of function words relies more on the target context, one may wonder to what extent attention is needed for preserving function words? To answer this question, we completely zero out the attention using ZeroOut. Row 5 shows that only 9% of function words were preserved in this method.

Moreover it can be seen that the model could not preserve any content word when this method is employed. Interestingly, we found that the preserved function words in this method were all occurrences of “,”. Apparently the decoder’s language model is so strong in predicting “,” without attention. This finding suggests that a basic representation of the source context is still necessary to preserve function words.

5.3 Highlighting top preserved tokens

An important question that may arise is whether each attention method tends to preserve a specific group of words. To address this question, we listed the top preserved function words and content words for all the proposed methods. We observed that they mostly preserve the same group of words but with different percentages. As a result, we only list the top preserved tokens for the aggregate method.

Table 3 contains the top 20 content words sorted by the number of times they were preserved. It is interesting to note that for many of these frequent tokens, more than half of their total occurrences are preserved without focusing on their corresponding translation in the source sentence (e.g., “going”, “know”, “thing”, etc).

In Table 4, we sort such tokens based on their coverage, which is the percentage of their total

Token	# preserved	Coverage
going	310	70%
people	237	46%
know	219	62%
world	215	67%
like	189	47%
think	176	50%
way	162	68%
get	160	53%
thing	147	79%
things	142	56%
time	139	54%
see	137	51%
years	136	64%
make	126	49%
little	113	55%
just	109	29%
really	93	37%
bit	92	88%
said	89	59%
got	86	59%

Table 3: Top 20 content words preserved by the aggregate method sorted by the number of times they were preserved.

Token	Coverage	Total
bit	88%	105
course	87%	91
thank	83%	89
thing	79%	186
fact	78%	74
half	78%	27
own	75%	75
ones	73%	30
states	73%	30
difference	71%	21
going	70%	444
turns	69%	26
way	68%	237
able	67%	85
world	67%	323
doing	66%	103
planet	65%	37
years	64%	212
know	62%	353
united	62%	21

Table 4: Top 20 content words preserved by the aggregate method sorted by percentage of their total occurrences that are preserved (*coverage*).

occurrences that are not affected when a counterfactual attention is applied⁴. We repeat the same process for function words (Table 5 and Table 6).

As evident from Table 5, we have successfully yielded the same token in 94% of the occurrences of the EOS token but with a contradictory explanation. This can be explained by the previous findings suggesting special hidden units keep track of translation length (Shi et al., 2016). As a result, EOS token is generated upon receiving signal from these units rather than using attention. This indicates that attention weights are highly unreliable for explaining the generation of EOS tokens. This is worth noting because early generation of the EOS token is often a major reason of the under-translation problem in NMT (Kuang et al., 2018). Thus, attention weights should not be used to debug early generation of EOS, and that some other underlying influence in the network (Ding et al., 2017) might be responsible for the model’s decision in this case.

5.4 Last encoder hidden state is a poor representation of the source context in the attentional model

The last encoder state in a non-attentional model is passed to the decoder as the representation of the source context (Cho et al., 2014) although the accuracy of this representation has been shown to degrade as sentence length increases (Bahdanau et al., 2014). We experiment with the `LastEncoderState` method to investigate how well the last encoder state can be representative of the source context in the attentional setting, and if exclusively attending to it can be used as a counterfactual attention.

Row 6 in Table 2 shows that there is a significant gap between the `LastEncoderState` and the counterfactual methods proposed in Table 2. A possible explanation for this result is that in the presence of attention mechanism, the model is trained to distribute the source-side information among the encoder states to be easier to select the relevant parts of the information with respect to the decoding state. Consequently, the last encoder state does no longer capture the whole information.

⁴We consider only the tokens that have appeared more than 20 times. The reason is that there are many preserved words that have appeared only once (coverage=1) and it is not clear if the coverage remains the same when frequency increases.

Token	# preserved	Coverage
,	7329	85%
EOS	6364	94%
the	5210	82%
.	3947	60%
of	3003	87%
to	2923	86%
and	2639	67%
a	2187	65%
that	1936	69%
i	1737	76%
's	1732	95%
you	1501	72%
it	1497	72%
is	1496	88%
in	1364	64%
we	1246	64%
they	624	69%
"	620	81%
have	613	70%
be	582	91%
't	580	96%
're	542	86%
this	541	42%
so	531	57%
are	526	77%
was	514	66%
do	433	77%
about	417	65%
what	415	61%
can	400	54%

Table 5: Top 30 function words preserved by the aggregate method sorted by the number of times they were preserved.

If we look at Table 7, we can see that the most covered functions words, are the words that usually appear at the end of the sentence (e.g., “EOS”, “.”, “?”, “!”). This is because most of the context captured by the last encoder state is centered around the last part of the sentence in which these tokens appear.

5.5 Attention of non-maximum source hidden states

Row 7 shows that a majority of output tokens can be preserved even when the model attends to a single source hidden state. Row 8 shows that when other source hidden states are uniformly combined, although the ratio of unaffected content words has increased by 3%, ratio of unaffected

Token	Coverage	Total
't	96%	602
's	95%	1819
EOS	94%	6748
be	91%	641
is	88%	1707
of	87%	3450
to	86%	3383
're	86%	631
,	85%	8582
'm	84%	311
been	82%	233
lot	82%	148
the	82%	6386
"	81%	770
are	77%	679
do	77%	565
i	76%	2290
who	73%	300
it	72%	2089
you	72%	2099
have	70%	876
up	70%	235
they	69%	904
that	69%	2812
well	67%	153
and	67%	3922
was	66%	774
were	65%	240
same	65%	154
a	65%	3369

Table 6: Top 30 function words preserved by the aggregate method sorted by coverage.

function words has increased by 15%. This again underlines the importance of the basic representation of the source context for generation of function words.

5.6 Summary

Our findings can be summarized as follows:

- It is possible to generate 68% of function words and 21% of content words with a counterfactual attention indicating unreliability of using attention weights as explanation.
- The generation of function words relies more on the target context, whereas the generation of content words relies more on the source context. This results in a higher likelihood of

Token	Coverage	Total
EOS	99%	6748
.	98%	6526
?	88%	589
't	86%	602
!	68%	22
"	60%	770
's	24%	1819
;	24%	63
are	18%	679
is	18%	1707

Table 7: Top 10 unaffected function words in the LastEncoderState method.

generation of preserved function words compared to that of preserved content words.

- Generation of EOS tokens cannot be reliably explained by using attention weights. Instead, this depends on the length of the target translation which is implicitly pursued by special hidden units. As a result, EOS token is emitted upon receiving a signal from these units rather than information from attention.
- The last encoder state is a poor representation of the source sentence and cannot be effectively used as the source context vector.
- It is possible to generate 86% of tokens by only using the source hidden state with the maximum attention and using other source hidden states uniformly suggesting that it may not be necessary to assign highly tuned weights to each source hidden state.

6 Related Work

Relevance-based interpretation is a common technique in analyzing predictions in neural models. In this method, inputs of a predictor are assigned a scalar value quantifying the importance of that particular input on the final decision. Saliency methods use the gradient of the inputs to define importance (Li et al., 2016a; Ghaeini et al., 2018; Ding et al., 2019). Layer-wise relevance propagation that assigns relevance to neurons based on their contribution to activation of higher-layer neurons is also investigated in NLP (Arras et al., 2016; Ding et al., 2017; Arras et al., 2017). Another method to measure relevance is by removing the input, and tracking the difference in in network’s

output (Li et al., 2016b). While these methods focus on explaining a model’s decision, Shi et al. (2016); Kádár et al. (2017); Calvillo and Crocker (2018) investigate how a particular concept is represented in the network.

Analyzing and interpreting the attention mechanism in NLP (Koehn and Knowles, 2017; Ghader and Monz, 2017; Tang and Nivre, 2018; Clark et al., 2019; Vig and Belinkov, 2019) is another direction that has drawn major interest. Although attention weights have been implicitly or explicitly used to explain a model’s decisions, the reliability of this approach is not proven. Several attempts have been made to investigate the reliability of this approach for explaining a models’ decision in NLP (Serrano and Smith, 2019; Baan et al., 2019; Jain et al., 2019; Jain and Wallace, 2019), and also in information retrieval (Jain and Madhyastha, 2019).

Our work was inspired by Jain and Wallace (2019). However, in this work we have focused on similar issues in neural machine translation which is has different challenges compared to text classification in terms of objective and architecture. Moreover, our paper studies the effect of different counterfactual attention methods.

7 Conclusion

Using attention weights to justify a model’s prediction is tempting and seems intuitive at the first glance. It is, however, not clear whether attention can be employed for such purposes. There might exist alternative attention weights resulting in the same decision by the model but promoting different contradictory explanation.

We propose several attention methods to create counterfactual attention weights from the original weights, and we measure to what extent these new contradictory weights can yield the same output as the original one. We find that in many cases, an output token can be generated even though a counterfactual attention is fed to the decoder. This implies that using attention weights to rationalize a model’s decision is not a reliable approach.

8 Future Work

In the future, we intend to study the extent to which attention weights correlate with importance measured by gradient-based methods. While we have separated function and content words in this work, we would like to extend our findings to other

categories such as parts of speech (POS) or out-of-vocabulary (OOV) words. Another logical investigation for future would be to address interpretability of copy mechanism in NMT (Gu et al., 2016). Proving the correlation between attention and the model predictions in more sophisticated attention models such as Transformer (Vaswani et al., 2017) is also worth exploring.

Acknowledgments

We would like to thank the anonymous reviewers for their helpful comments. The research was also partially supported by the Natural Sciences and Engineering Research Council of Canada grants NSERC RGPIN-2018-06437 and RGPAS-2018-522574 and a Department of National Defence (DND) and NSERC grant DGDND-2018-00025 to the second author.

References

- Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2016. [Explaining predictions of non-linear classifiers in NLP](#). In [Proceedings of the 1st Workshop on Representation Learning for NLP](#), pages 1–7, Berlin, Germany. Association for Computational Linguistics.
- Leila Arras, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2017. [Explaining recurrent neural network predictions in sentiment analysis](#). In [Proceedings of the 8th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis](#), pages 159–168, Copenhagen, Denmark. Association for Computational Linguistics.
- Joris Baan, Maartje ter Hoeve, Marlies van der Wees, Anne Schuth, and Maarten de Rijke. 2019. [Do transformer attention heads provide transparency in abstractive summarization?](#) [arXiv preprint arXiv:1907.00570](#).
- Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. [On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation](#). [PLOS ONE](#), 10(7):1–46.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. [Neural machine translation by jointly learning to align and translate](#). [arXiv preprint arXiv:1409.0473](#).
- Jesús Calvillo and Matthew Crocker. 2018. [Language production dynamics with recurrent neural networks](#). In [Proceedings of the Eight Workshop on Cognitive Aspects of Computational Language Learning and Processing](#), pages 17–26, Melbourne. Association for Computational Linguistics.

- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In [Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing \(EMNLP\)](#), pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. [What does BERT look at? an analysis of BERT’s attention](#). In [Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP](#), pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Shuoyang Ding, Hainan Xu, and Philipp Koehn. 2019. [Saliency-driven word alignment interpretation for neural machine translation](#). In [Proceedings of the Fourth Conference on Machine Translation](#), pages 1–12, Florence, Italy. Association for Computational Linguistics.
- Yanzhuo Ding, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. [Visualizing and understanding neural machine translation](#). In [Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 1150–1159, Vancouver, Canada. Association for Computational Linguistics.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. [Pathologies of neural models make interpretations difficult](#). In [Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing](#), pages 3719–3728, Brussels, Belgium. Association for Computational Linguistics.
- Hamidreza Ghader and Christof Monz. 2017. [What does attention in neural machine translation pay attention to?](#) In [Proceedings of the Eighth International Joint Conference on Natural Language Processing \(Volume 1: Long Papers\)](#), pages 30–39, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Reza Ghaeini, Xiaoli Fern, and Prasad Tadepalli. 2018. [Interpreting recurrent and attention-based neural models: a case study on natural language inference](#). In [Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing](#), pages 4952–4957, Brussels, Belgium. Association for Computational Linguistics.
- Amirata Ghorbani, Abubakar Abid, and James Zou. 2019. [Interpretation of neural networks is fragile](#). In [Proceedings of the AAAI Conference on Artificial Intelligence](#), volume 33, pages 3681–3688.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In [Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 1631–1640, Berlin, Germany. Association for Computational Linguistics.
- Rishabh Jain and Pranava Madhyastha. 2019. [Model explanations under calibration](#). [arXiv preprint arXiv:1906.07622](#).
- Sarthak Jain, Ramin Mohammadi, and Byron C. Wallace. 2019. [An analysis of attention over clinical notes for predictive tasks](#). In [Proceedings of the 2nd Clinical Natural Language Processing Workshop](#), pages 15–21, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not explanation](#). 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 3543–3556, Minneapolis, Minnesota. Association for Computational Linguistics.
- Akos Kádár, Grzegorz Chrupała, and Afra Alishahi. 2017. [Representation of linguistic form and function in recurrent neural networks](#). [Computational Linguistics](#), 43(4):761–780.
- Andrej Karpathy, Justin Johnson, and Li Fei-Fei. 2015. [Visualizing and understanding recurrent networks](#). [arXiv preprint arXiv:1506.02078](#).
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). In [Proceedings of ACL 2017, System Demonstrations](#), pages 67–72, Vancouver, Canada. Association for Computational Linguistics.
- 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 Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions](#), pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- 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.
- Shaohui Kuang, Junhui Li, António Branco, Weihua Luo, and Deyi Xiong. 2018. [Attention focusing for neural machine translation by bridging source and target embeddings](#). In [Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics \(Volume 1: Long Papers\)](#), pages 1767–1776, Melbourne, Australia. Association for Computational Linguistics.

- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2016a. [Visualizing and understanding neural models in NLP](#). In [Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies](#), pages 681–691, San Diego, California. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016b. [Understanding neural networks through representation erasure](#). [arXiv preprint arXiv:1612.08220](#).
- Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. [Effective approaches to attention-based neural machine translation](#). In [Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing](#), pages 1412–1421, Lisbon, Portugal. Association for Computational Linguistics.
- Grégoire Montavon, Sebastian Lapuschkin, Alexander Binder, Wojciech Samek, and Klaus-Robert Müller. 2017. [Explaining nonlinear classification decisions with deep taylor decomposition](#). [Pattern Recognition](#), 65:211–222.
- Sofia Serrano and Noah A. Smith. 2019. [Is attention interpretable?](#) In [Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics](#), pages 2931–2951, Florence, Italy. Association for Computational Linguistics.
- Xing Shi, Kevin Knight, and Deniz Yuret. 2016. [Why neural translations are the right length](#). In [Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing](#), pages 2278–2282, Austin, Texas. Association for Computational Linguistics.
- Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. 2013. [Deep inside convolutional networks: Visualising image classification models and saliency maps](#). [arXiv preprint arXiv:1312.6034](#).
- Hendrik Strobelt, Sebastian Gehrmann, Michael Behrisch, Adam Perer, Hanspeter Pfister, and Alexander M Rush. 2018. [Seq2seq-v is: A visual debugging tool for sequence-to-sequence models](#). [IEEE transactions on visualization and computer graphics](#), 25(1):353–363.
- Hendrik Strobelt, Sebastian Gehrmann, Hanspeter Pfister, and Alexander M Rush. 2017. [Lstmvis: A tool for visual analysis of hidden state dynamics in recurrent neural networks](#). [IEEE transactions on visualization and computer graphics](#), 24(1):667–676.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. [Sequence to sequence learning with neural networks](#). In [Advances in neural information processing systems](#), pages 3104–3112.
- Rico Tang, Gongbo Sennrich and Joakim Nivre. 2018. [An analysis of attention mechanisms: The case of word sense disambiguation in neural machine translation](#). In [Proceedings of the Third Conference on Machine Translation: Research Papers](#), pages 26–35, Belgium, Brussels. Association for Computational Linguistics.
- Zhaopeng Tu, Yang Liu, Zhengdong Lu, Xiaohua Liu, and Hang Li. 2017. [Context gates for neural machine translation](#). [Transactions of the Association for Computational Linguistics](#), 5:87–99.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In [Advances in neural information processing systems](#), pages 5998–6008.
- Jesse Vig and Yonatan Belinkov. 2019. [Analyzing the structure of attention in a transformer language model](#). In [Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP](#), pages 63–76, Florence, Italy. Association for Computational Linguistics.
- Matthew D Zeiler and Rob Fergus. 2014. [Visualizing and understanding convolutional networks](#). In [European conference on computer vision](#), pages 818–833. Springer.