



HAL
open science

A Study of the Plausibility of Attention between RNN Encoders in Natural Language Inference

Duc Hau Nguyen, Guillaume Gravier, Pascale Sébillot

► **To cite this version:**

Duc Hau Nguyen, Guillaume Gravier, Pascale Sébillot. A Study of the Plausibility of Attention between RNN Encoders in Natural Language Inference. ICMLA 2021 - 20th IEEE International Conference on Machine Learning and Applications, Dec 2021, Pasadena, United States. pp.1-7. hal-03372669

HAL Id: hal-03372669

<https://hal.science/hal-03372669>

Submitted on 11 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Study of the Plausibility of Attention between RNN Encoders in Natural Language Inference

Duc Hau Nguyen
IRISA, CNRS, INSA Rennes
Rennes, France
duc-hau.nguyen@irisa.fr

Guillaume Gravier
IRISA, CNRS
Rennes, France
guillaume.gravier@irisa.fr

Pascale Sébillot
IRISA, INSA Rennes
Rennes, France
pascale.sebillot@irisa.fr

Abstract—Attention maps in neural models for NLP are appealing to explain the decision made by a model, hopefully emphasizing words that justify the decision. While many empirical studies hint that attention maps can provide such justification from the analysis of sound examples, only a few assess the plausibility of explanations based on attention maps, i.e., the usefulness of attention maps for humans to understand the decision. These studies furthermore focus on text classification. In this paper, we report on a preliminary assessment of attention maps in a sentence comparison task, namely natural language inference. We compare the cross-attention weights between two RNN encoders with human-based and heuristic-based annotations on the eSNLI corpus. We show that the heuristic reasonably correlates with human annotations and can thus facilitate evaluation of plausible explanations in sentence comparison tasks. Raw attention weights however remain only loosely related to a plausible explanation.

Index Terms—attention maps, explainability, natural language inference

I. INTRODUCTION

Attention mechanisms in deep learning have gained huge popularity in the past 5 years due to the benefits on performance across many tasks. In natural language processing (NLP) in particular, these mechanisms are now widely used, either as an attention layer on top of a recurrent network [1]–[3] or directly as a self-attention model [4]–[6].

In addition to improved performance, attention weights have also attracted interest as a way to help human understand how deep learning models process data and make decisions [7]–[9]. Typically, in NLP, an attention mechanism provides a weight on each input word, presumably proportional to the word’s relevance regarding the decision, and weights can be visualized as a heat map, known as an *attention map*.

However, qualifying the usefulness for a human to understand the decision that was made by a model—a concept globally known as *explainability*—remains a challenge. The challenge arises from two concerns: on the one hand, explainability isn’t directly related to performance on the end task; on the other hand, the cost of post-hoc user studies or the lack of annotated data for attention layers impose severe limits, as attention maps are usually one of several layers within a complex architecture. Moreover, regardless of performance, most of the evaluations of attention maps for explainability in NLP were done empirically at the sentence level, with models

that were primarily designed for performance rather than for explainability.

It is now well-known that a model with a good performance on a certain task does not guarantee a good explanation, which makes the end-task performance a poor proxy to assess explainability. Recent works, such as [10]–[12], disproved the uniqueness of explanation for the same prediction by showing that there exists an adversarial model maintaining similar prediction but having a completely different attention map.

Evaluation of the usefulness of attention maps regarding explainability must therefore focus on human assessment of the attention map itself. Two natural approaches can be considered to this end: posterior user studies of attention maps, which are costly and hardly provide sound criteria to train attention mechanisms with explainability in mind; prior annotation of important input token/words that should be emphasized to explain the decision.

In both cases, it is nevertheless unclear whether human judgement necessarily provides a good baseline for explanation, depending on what is meant by explanation. Interestingly, the notions of faithfulness and plausibility were recently introduced [13]. Briefly speaking, faithfulness reflects how much an attention map “reflects the model’s reasoning process”—in other words, how a model makes the decision with the same chosen words or a chosen part of input in general—, while plausibility reflects how useful the attention map is to humans.

Most recent studies on explainability derived from attention maps assess the faithfulness of the attention mechanism, e.g., qualitatively looking at attention weights on a few examples [9], [14]–[16]. Addressing plausibility is however much less studied, because of the lack of ground truth and the difficulty to define one. In one of the few studies on plausibility of attention maps in NLP [17], annotators were asked to classify attention maps as informative or not in a medical text classification task.

In this paper, we investigate the factors that influence the plausibility of an attention mechanism between two LSTM recurrent encoders on a natural language inference (NLI) task where two sentences are compared to detect entailment or contradiction. The choice of the task responds to two criteria apart from novelty: it is significantly more complex than text classification, as considered in [17]; it is close to an information retrieval task where links between text fragments must be established. This last consideration is important to

us as learning to link two sentences is critical in many applications, often requires plausible explanations for users to understand why a link is proposed, and hardly comes with a ground truth explanation from which an explanation model can be learned.

Taking NLI as an emblematic task where two sentences have to be compared, we take advantage of human annotations within the eSNLI dataset [18], an extension of the standard Stanford Natural Language Inference dataset [19] that comes with human annotations of the words that should be highlighted in the premise and hypothesis sentences. We compare attention weights between two RNN encoders with human annotations and with a simple heuristic, close to the one proposed in [17] for text classification. Although based on a different criterion than the annotation instructions in eSNLI, the heuristic is fully justified for the entailment task, in particular to get plausible explanations in the case of entailment. We also experimentally show that the heuristic reasonably correlates with human annotations and can thus facilitate evaluation of plausible explanations in sentence comparison tasks.

II. RELATED WORKS

Originally, attention mechanisms were introduced in machine translation to help the focus on contextual words in sequence encoding [1], [2]. A brief visualization from the original papers demonstrating responsible tokens for each translation words has inspired other studies to use attention maps [9], [16] to demonstrate influential parts of each instance’s decision (also referred as local explanation [9], [20]). Similar approaches were taken for self-attention models.

Following this trend, [10] questioned the validity of the attention map and resulted in disproving the faithfulness in explanation in some tasks. In other words, we can have a different attention map to get the same prediction result. [12] has disproved that an attention map does not have the linear contribution to the final prediction, as it can mask out important words and still get the approximate prediction. To explain this phenomenon, [11] has introduced the concept of faithfulness when trying to evaluate an attention mechanism, which refers to the capability that an attention map can reflect the model reasoning. The authors argued that the attention mechanism does reflect the model’s deduction process but is not necessarily comprehensible to human. Another factor is the contribution of the attention layer in the task: if this layer is replaceable, then it is unlikely to have any explanation property. [21] has bolstered this argument, demonstrating that in single sequence inference tasks (sentiment or topic classifications for example), the attention score becomes just a gate unit used in a convolution model [22].

While the faithfulness of the model has been the subject of several works, the plausibility of attention maps is still rarely studied to our knowledge. While there is a direct way to measure plausibility of attention map in computer vision-related tasks, as in [23] for example, most of NLP corpus does not have a direct way to evaluate against human-annotated

ground truth. Instead, [24] studies the usefulness of attention map in a text classification task, asking workers to rank how much he or she is confident to find the correct label given the attention map. Others, such as [17], tried to evaluate the plausibility of their heuristic map, asking physicians to rank it against three other heuristics. The study here again relies on text classification tasks. Despite the claim of a satisfactory explainability of their model, result showed that a heuristic map based on similarity between words is more informative than the attention map is.

In this paper, we therefore extend the study of the plausibility of attention maps to a sentence comparison task—natural language inference, not restricting ourselves to a specific domain and taking advantage of existing human annotations in the eSNLI dataset.

III. REFERENCE ANNOTATIONS IN NLI

Natural language inference, also known as recognizing textual entailment, considers two sentences, namely a premise and an hypothesis, and consists in determining whether the hypothesis is entailed by or contradicts the premise, or if there is neither entailment nor contradiction. This typically results in a three class classification problem.

In this work, we employ the eSNLI corpus, an extension of the standard SNLI dataset where relevant words for a plausible explanation of the relation between the two sentences are highlighted by annotators. Annotators were asked to “focus on the non-obvious elements that induce the given relation, and not the parts of the premise that are repeated identically in the hypothesis”—see [18], Sec. 3 for details. Presumably, the words highlighted by the annotators, denoted as highlight map, are thus the ones we would like a plausible attention-based explanation model to emphasize. Highlight maps are provided for the entailment and contradiction classes. We focus mostly on the former in this study, sticking to the information retrieval scenario described in the introduction.

The eSNLI dataset provides approx. 180,000 pairs of sentences for each class in the train set, and approx. 3,000 pairs per class in the dev and test sets respectively. In these two last sets, the rate of out-of-vocabulary words (i.e., words not appearing in the train set) is around 5.5% (resp. 6%) without (resp. with) lemmatization. We report in Table I a brief analysis of the highlighted words in terms of their part-of-speech (POS) category. Note that a large part of highlighted words are among verb, noun and adjective, which suggests that certain grammatical categories offer better plausibility. Overall, 18% of the words were highlighted by the annotators.

Finally, let us note that the SNLI corpus is known to have artifacts, where some lexical fields appear mostly in one class. As reported in [25], words related to animals and outdoor are much more frequent in entailment than in the other two classes and a simple model extracting terms in the hypothesis can predict well the class. Far from being an issue, this bias is beneficial to us as it enables to distinguish between faithfulness (how much the attention map explains the model’s reasoning) and plausibility (how much the attention map is useful for

TABLE I: Statistics on the POS category of highlighted words.

		<i>train</i>	<i>dev</i>	<i>test</i>
% POS tag	VERB	21.90%	21.20%	20.8%
	NOUN	49.18%	48.4%	43.92%
	ADJ	9.08%	9.14%	7.98%
	NUM	2.89%	1.94%	1.36%
	ADP	0.15%	2.31%	7.55%
	DET	15.52%	9.34%	8.30%
% (VERB+NOUN+ADJ)		73.42%	78.75%	72.70%
% of highlight words		18.00%	18.21%	18.25%

a human to interpret the decision). Clearly, if attention maps focus on the model’s reasoning in this bias context, they should be poorly plausible unless one’s aware of the bias in the data.

IV. A HEURISTIC-BASED PLAUSIBLE ATTENTION MAP

Apart from human annotations in the eSNLI corpus, we designed a heuristic-based approach to plausible attention maps inspired from [26]. The rationale for using a heuristic that can be computed from the input sentences is that if valid, i.e., correlated with human annotations, it might serve as a criterion to drive the model’s attention weights towards a plausible explanation. As of now, the heuristic is used for evaluation purposes.

The heuristic that we propose focus on the entailment class where the premise and hypothesis sentences are strongly related (i.e., we can assert that hypothesis is true given the premise) and seek to highlight words that are close in meaning between the two sentences. Note that the same heuristic might be justified for the contradiction class, however not for the neutral one. In the case of entailment, it seems a reasonable assumption to state that plausibility of the explanation is related to words with close meanings between the two sentences, in particular verbs, nouns and adjectives.

Formally, let w_i , $i \in [1, n]$ denote a word token in one text sequence of length n and v_j , $j \in [1, m]$ a word token in the other text sequence of length m . The intuition is that we measure how relevant each word w_i is by comparing it to every word v_j in the other sentence. We note a function *similarity* that scores the similarity between two words and sum over the similarity with respect to w_i to get its relative strength among other words in the sequence. A ”good” word should have a better score, since it has short distant to many words in the other side.

A heuristic score $h(w_i)$ is defined by applying a normalization to reduce the magnitude of w_i in long sentences. Since words may have more relatives in longer sentence, the sum tends to be higher. Denoting as σ the normalizing function, the heuristic is given as

$$h(w_i) = \sigma \left(\sum_{v_j \notin S} \text{similarity}(w_i, v_j) \right), \quad (1)$$

where S is the set of stop words. We used the cosine similarity between embeddings of the words as similarity function. The

normalizing function σ is a rescaling function in the range $[0, 1]$, i.e.,

$$\sigma(x) = \frac{x - \min(x)}{\max(x) - \min(x)}. \quad (2)$$

A remark is that the heuristic is similar to the ”word by word” attention [26], but we rather use the similarity dot function as in [1]. Fig. 1 illustrates the heuristic attention map on an example of a sentence pair from the corpus: each entry in the attention map is set to $\text{similarity}(w_i, v_j) \in [0, 1]$, a high value indicating a relevant word. Stop words are masked out.

V. MODEL-BASED ATTENTION MAP

The study relies on a fairly standard architecture for NLI with two LSTM encoders (contextualization layer) linked through a cross-attention model and a decision layer as illustrated in Fig. 2. We make the assumption that the premise and hypothesis sentences have the same statistical attributes: they are thus processed with the same embedding and contextualization layer (i.e., shared weights). The sentence embedding resulting from the contextualization layer could have been used for classification, but we add an attention layer not only to (a) improve the model decision in the classifying layer, but also

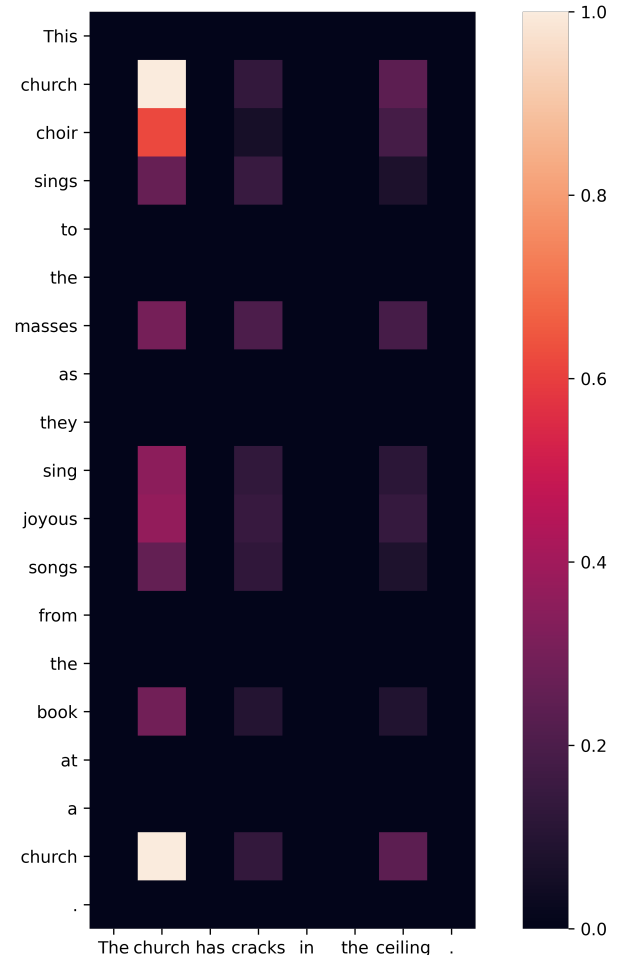


Fig. 1: Example of a heuristic-based plausible attention map.

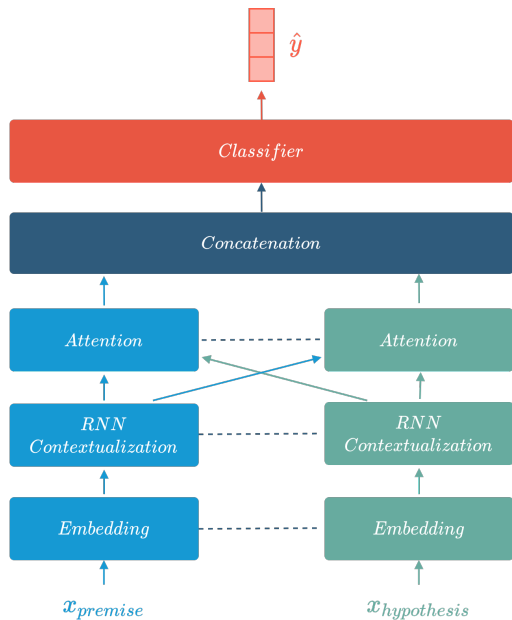


Fig. 2: Overall model architecture. The dash lines indicate that layers share the weights; in other words, we use the same first 3 layers (*Embedding*, *Contextualization*, *Attention*) for both *premise* and *hypothesis*.

to (b) benefit from the fact that the attention weight in this position can tell us which elements are taken into consideration by the classifier.

More precisely, each word in the sentences, after being transformed by a tokenizer, is represented by an embedding vector w_i (and v_i for word vector) through the first layer, i denoting the token position. The embedding layer weight is initialized from pretrained word vectors issued by the GloVe algorithm but we keep the weight trainable. A word embedding vector might encode ambiguous information if the context is disregarded. The representation is therefore refined by the context (i.e., the surrounding words) with a LSTM to contextualize word vectors into hidden vectors, denoted $h = [h_1, h_2, \dots, h_m]$ for a sentence of m word tokens, according to

$$h_t = \text{LSTM}(w_t, h_{t-1}) . \quad (3)$$

We chose to initialize the hidden vectors as null vectors $h_0 = \vec{0}$. The last hidden vector h_m is known to capture the sentence meaning and is often referred to as the sentence embedding. If we denote \bar{h}_v the hidden sequence for the remaining text sequence, we can define an attention mechanism that compares each word h_i of one sentence (the premise or the hypothesis) with the sentence embedding \bar{h}_n of the opposite sentence. As in the heuristic, the *cosine* similarity is used to estimate the proximity between h_i and \bar{h}_n . Given the score for each word, the formal attention weight vector $\alpha = [\alpha_1, \dots, \alpha_m]$ is a softmax normalized vector to yield a relative magnitude, the definition of the attention weight at each word position i being

defined as

$$\alpha_i = \frac{\exp(h_i^\top \bar{h}_n)}{\sum_{h_k \in h} \exp(h_k^\top \bar{h}_n)} . \quad (4)$$

In addition, the attention mechanism produces a new enhanced sentence embedding c , also referred to as the context vector in [1]. This embedding is then concatenated and blended with the RNN embedding h_n by a fully connected layer. The two new merged embeddings are then concatenated for classification. The attention process is illustrated in Fig. 3.

For the classifier part, several fully connected layers are stacked, with ReLU and tanh experimented as activation. A final softmax is applied at the final layer to classify into 3 classes, with cross-entropy as the loss function in training.

This model is inspired by the decoder-based architecture from [1]. We use the last final layer from the hidden representation of a BiLSTM to perform attention against each word in the opposite sequence. As shown in Fig. 3, the dot product is chosen as the attention score, as we need to compute the semantically similar words between the text sequences regarding the task. Table II reports the optimal model hyper parameters by maximizing the F-1 score on the dev set of eSNLI.

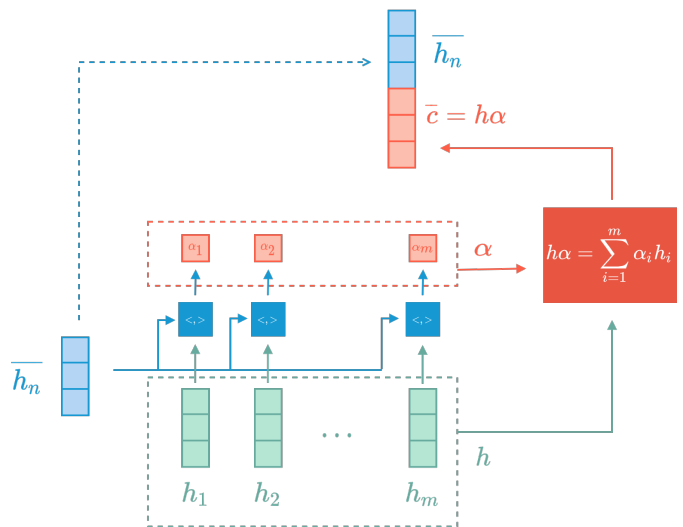


Fig. 3: The attention mechanism used in the experimentation. The intuition is that α vector should give how much relevant a word is comparing to the sentence embedding of the other side \bar{h}_m .

TABLE II: Optimal model configuration.

Configuration	optimal setting
Embedding	pretrained GloVe, 300d
LSTM	BiLSTM, 1 layer, 300d
MLP Classifier	1 hidden layer, <i>relu</i> activation

VI. EXPERIMENTAL EVALUATION

A. Experiment Setup

The model is trained on the entire train set of eSNLI with hyperparameters optimized on the dev set. Attention maps are evaluated on the samples from the entailment class from the test set as the idea is to provide explainability on the positive class, i.e., when something of interest is detected.

B. Results

The first question that we address is whether the model-based map matches the heuristic map or the human-annotated map. Let us first quantitatively analyze the proximity against human annotation.

To compare against the binary human annotation, the attention values in the attention maps, whether model-based or heuristic-based, must be turned into binary values by selecting a threshold $\epsilon \in [0, 1]$, the attention values being 1 above ϵ and 0 otherwise. By investigating the true positive and false positive rates as a function of ϵ , we obtain a ROC curve as shown in Fig. 4. A good matching pair has a high true positive rate in almost all ϵ values so the corresponding ROC curve leans towards the upper left corner of the graphic. Results in Fig. 4 show a good match between the heuristic map and human annotations.

Can we conclude the model-based map is closer to the heuristic one? Let us establish a similar curve, this time we comparing the model-based and annotated maps against the heuristic map—or, similarly, we use the heuristic map as ground truth. However, turning a heuristic map into a ground truth label needs to define a $\epsilon_{\text{heuristic}} \in [0, 1]$ as previously. The problem is that for each $\epsilon_{\text{heuristic}}$, there is a different shape of the ROC curves pairs. We thus use the area under the curve (AUC) to measure how relatively high a curve is—the higher the curve, the better the map matches its ground truth—, reporting AUC for different values of $\epsilon_{\text{heuristic}}$.

Table III reports the AUC values for some values of $\epsilon_{\text{heuristic}}$. We observe that with $\epsilon_{\text{heuristic}} \in [0.42, 0.5]$, the model attention matches better the heuristic, while overall, the heuristic matches better the annotation. One reason is that at low $\epsilon_{\text{heuristic}}$ values, the heuristic has much more highlights across sentences that cover even parts non highlighted by humans. Meanwhile, model attention tends to distribute highlights on many unimportant tokens. Figure 5 also shows that for small values of $\epsilon_{\text{heuristic}}$, the model attention matches better the heuristic than the human highlight map. For 85.24% of $\epsilon_{\text{heuristic}}$ values, the human annotation matches the heuristic map better than the model attention. We can now answer the first question: the attention map has not enough plausibility compared to neither human nor heuristic ground truth.

The next interesting question is whether the heuristic map reflects how a human picks up words to explain the label decision. In other words, does the distribution of attention values (actual values, not thresholded binary values) match between the model and the human annotations. For the latter, binary values are normalized so that attention values sum to

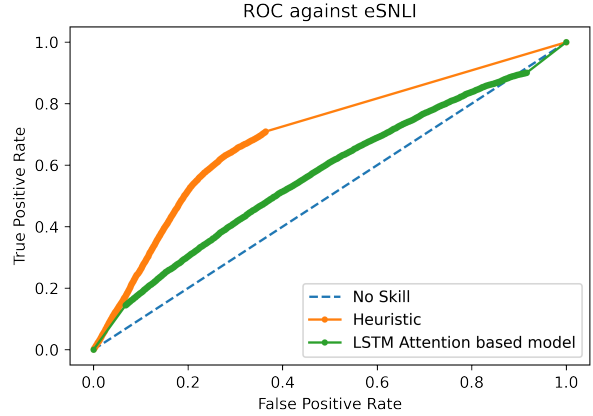


Fig. 4: ROC curve that measures how much a model matches the ground truth. The orange (resp. green) line compares the heuristic (resp. model-based) map with the human annotation. The dash blue line in diagonal is the worst prediction, where a system highlights every words.

TABLE III: AUC as a function of $\epsilon_{\text{heuristic}}$.

$\epsilon_{\text{heuristic}}$	AUC_{human}	AUC_{model}
0.4	0.380450	0.342720
0.5	0.634234	0.522256
0.6	0.635616	0.527185
0.7	0.648283	0.557091
0.8	0.645316	0.571382
0.9	0.627151	0.555492

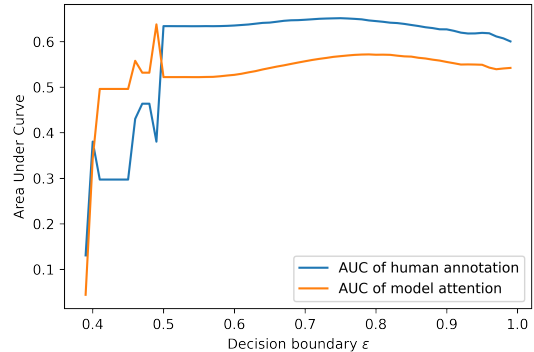


Fig. 5: AUC of human annotation map (orange) and model attention (blue) as a function of $\epsilon_{\text{heuristic}}$.

unity in each sentence. On a quantitative global scale, we measure the correlation between the distribution of attention values between the two maps in Table IV. The small values of the p-value for both Pearson and Spearman correlation reject the null hypothesis H_0 that the correlation is statistically insignificant, i.e., the correlation values show a low correlation between the two maps.

However, the heuristic map at a per instance scale still has a dispersion across sentence pairs as shown in Fig. 6

TABLE IV: Correlation between heuristic and human annotation.

	Pearson		Spearman	
	Correlation	p-value	Correlation	p-value
Premise	0.5227	$< 10^{-300}$	0.5337	$< 10^{-300}$
Hypothesis	0.5227	$< 10^{-300}$	0.4605	$< 10^{-300}$
General	0.5206	$< 10^{-300}$	0.4426	$< 10^{-300}$

where we report the distribution of several metrics measured at the sentence level rather than at the global level. We report three different metrics: the Jensen-Shannon divergence (JS-divergence) between the two distributions of attention values (the lower, the better), as well as the Spearman and Pearson correlations (the higher, the better). The distribution of the JS-divergence in Fig. 6 shows that most of the distribution significantly differ. Similarly, the high p-values in Spearman and Pearson correlation tests reject any correlation between the two distributions.

A qualitative investigation can confirm this last observation. We show examples of two sequence pairs of the entailment class in Fig. 7, where the heuristic can provide better coverage on human annotation highlights, while the model-based one mismatches many words. This result explains why, on a global scale, heuristics still have a relation with human annotation, though not a perfect match.

In summary, the attention between RNN encoders has a very low plausibility, even compared to the heuristic map, one reason is that the model focuses mostly on unimportant grammatical parts of the sentence (determinants, punctuation, etc.) and skip the informative syntactic groups for human (verb, noun, adjective). On the other hand, the heuristic is proved to be a good baseline to measure plausibility in NLI task. Confirming [27], the embedding presentation can provide useful information for word choice and may result in higher plausibility in the heuristic map in NLI task.

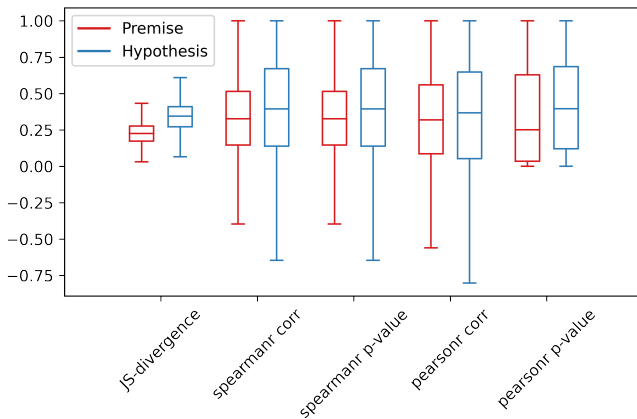


Fig. 6: Distribution of Spearman and Pearson correlation ratios across sentences (premise in red, hypothesis in blue).

VII. CONCLUSION

In this paper, we performed an empirical study on the plausibility of the attention mechanism in a sentence comparison task. The attention map is built on top of RNN-encoders—BiLSTM models. In this work, we exploited human explaining annotations from the eSNLI corpus. The explanations in the corpus are provided by an explaining phrase for the label and highlighted maps. We used the latter as a ground truth to evaluate the model attention map. In addition, we designed a heuristic method, based on words similarity and grammatical attributes, as an additional second ground-truth to assess the plausibility. The results showed that the model-based attention map is not close to both ground-truth maps. The qualitative analysis shows that the model gives scattered attention weights, most of them on stop words, which bring few indicative information to humans. However, the heuristic map, based on initial embedding words, shows a better coverage of these words.

This observation shows the importance of the word embedding vectors which could bring useful information to the attention mechanism to improve its plausibility. In future works, the heuristic could be used as a tool for evaluation in NLI tasks for corpus without annotations. As an extension, the heuristic could serve as a complementary to a regularization proposed in [27] to enhance the plausibility without sacrificing the faithfulness of the attention mechanism.

REFERENCES

- [1] T. Luong, H. Pham, and C. D. Manning, “Effective Approaches to Attention-based Neural Machine Translation,” in *Proc. of 2015 Conf. on Empirical Methods in Natural Language Processing*, 2015.
- [2] D. Bahdanau, K. Cho, and Y. Bengio, “Neural Machine Translation by Jointly Learning to Align and Translate,” in *Proc. of 3rd International Conference on Learning Representations*, 2015.
- [3] Q. Chen, X. Zhu, Z.-H. Ling, S. Wei, H. Jiang, and D. Inkpen, “Recurrent Neural Network-Based Sentence Encoder with Gated Attention for Natural Language Inference,” in *Proc. of 2nd Workshop on Evaluating Vector Space Representations for NLP*, 2017.
- [4] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is All you Need,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding,” in *Proc. of Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.
- [6] T. Shen, T. Zhou, G. Long, J. Jiang, S. Pan, and C. Zhang, “DiSAN: Directional Self-Attention Network for RNN/CNN-Free Language Understanding,” in *Proc. of AAAI Conference on Artificial Intelligence*, Apr. 2018.
- [7] S. Wang and J. Jiang, “Learning Natural Language Inference with LSTM,” in *Proc. of Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2016.
- [8] Y. Wang, M. Huang, X. Zhu, and L. Zhao, “Attention-based LSTM for Aspect-level Sentiment Classification,” in *Proc. of 2016 Conf. on Empirical Methods in Natural Language Processing*, 2016.
- [9] R. Ghaeini, X. Fern, and P. Tadepalli, “Interpreting Recurrent and Attention-Based Neural Models: a Case Study on Natural Language Inference,” in *Proc. of 2018 Conf. on Empirical Methods in Natural Language Processing*, 2018.
- [10] S. Jain and B. C. Wallace, “Attention is not Explanation,” in *Proc. of Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019.

	Premise	Hypothesis		Premise	Hypothesis
	Annotation map	Two men on bicycles competing in a race . People are riding bikes .	Annotation map	Two women having drinks and smoking cigarettes at the bar .	Two women are at a bar .
	Model-based map	Two men on bicycles competing in a race . People are riding bikes .	Model-based map	Two women having drinks and smoking cigarettes at the bar .	Two women are at a bar .
[Short pair]	Heuristic map	Two men on bicycles competing in a race . People are riding bikes .	Heuristic map	Two women having drinks and smoking cigarettes at the bar .	Two women are at a bar .

(b) Long pair

Fig. 7: Two examples comparing human-based, model-based and heuristic-based attention maps.

- [11] S. Wiegrefe and Y. Pinter, "Attention is not Explanation," in *Proc. of 2019 Conf. on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, 2019.
- [12] S. Serrano and N. A. Smith, "Is Attention Interpretable?" in *Proc. of 57th Annual Meeting of the Association for Computational Linguistics*, 2019.
- [13] A. Jacovi and Y. Goldberg, "Towards Faithfully Interpretable NLP Systems: How Should We Define and Evaluate Faithfulness?" in *Proc. of 58th Annual Meeting of the Association for Computational Linguistics*, 2020.
- [14] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention," in *Proc. of 32nd International Conference on Machine Learning*, F. Bach and D. Blei, Eds., 2015.
- [15] E. Choi, M. T. Bahadori, J. Sun, J. Kulas, A. Schuetz, and W. Stewart, "RETAIN: An Interpretable Predictive Model for Healthcare using Reverse Time Attention Mechanism," in *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [16] Z. Lin, M. Feng, C. N. d. Santos, M. Yu, B. Xiang, B. Zhou, and Y. Bengio, "A Structured Self-attentive Sentence Embedding," in *Proc. of 5th International Conference on Learning Representations*, 2017.
- [17] J. Mullenbach, S. Wiegrefe, J. Duke, J. Sun, and J. Eisenstein, "Explainable Prediction of Medical Codes from Clinical Text," in *Proc. of Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2018.
- [18] O.-M. Camburu, T. Rocktäschel, T. Lukasiewicz, and P. Blunsom, "e-SNLI: Natural Language Inference with Natural Language Explanations," in *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [19] S. R. Bowman, G. Angeli, C. Potts, and C. D. Manning, "A large annotated corpus for learning natural language inference," in *Proc. of 2015 Conference on Empirical Methods in Natural Language Processing*, 2015.
- [20] L. H. Gilpin, D. Bau, B. Z. Yuan, A. Bajwa, M. Specter, and L. Kagal, "Explaining Explanations: An Overview of Interpretability of Machine Learning," in *Proc. of 5th Int. Conf. on Data Science and Advanced Analytics*, 2018.
- [21] S. Vashishth, S. Upadhyay, G. S. Tomar, and M. Faruqui, "Attention Interpretability Across NLP Tasks," *arXiv:1909.11218*, 2019.
- [22] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language Modeling with Gated Convolutional Networks," in *Proc. of 34th International Conference on Machine Learning*, 2017.
- [23] J. Wu and R. Mooney, "Faithful Multimodal Explanation for Visual Question Answering," in *Proc. of the ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, 2019.
- [24] D. Nguyen, "Comparing Automatic and Human Evaluation of Local Explanations for Text Classification," in *Proc. of Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2018.
- [25] S. Gururangan, S. Swayamdipta, O. Levy, R. Schwartz, S. Bowman, and N. A. Smith, "Annotation Artifacts in Natural Language Inference Data," in *Proc. of Conf. of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2018.
- [26] T. Rocktäschel, E. Grefenstette, K. M. Hermann, T. Kočiský, and P. Blunsom, "Reasoning about Entailment with Neural Attention," in *Proc. of International Conference on Learning Representations*, 2016.
- [27] M. Tutek and J. Snajder, "Staying True to Your Word: (How) Can Attention Become Explanation?" in *Proc. of 5th Workshop on Representation Learning for NLP*, 2020.