Interpretability of Attention Mechanisms in a Portuguese-Based Question Answering System about the Blue Amazon

Stefano Spindola[1,†,[0000−0002−3604−9346]], **Marcos Menon José**[1,†,[0000−0003−4663−4386]], **André Seidel Oliveira**[1,[0000−0001−6551−6911]], **Flávio Nakasato Cação**[1,[0000−0003−4771−6009]], **Fábio Gagliardi Cozman**[1,[0000−0003−4077−4935]]

[1] Escola Politécnica, Universidade de São Paulo, Sao Paulo, Brazil
† These authors contributed equally to the work

{stefano.spindola, marcos.jose, andre.seidel, flavio.cacao, fgcozman}@usp.br

Abstract. *The Brazilian Exclusive Economic Zone, or the "Blue Amazon", with its extensive maritime area, is the primary means of transport for the country's foreign trade and is important due to its oil reserves, gas and other mineral resources, in addition to the significant influence on the Brazilian climate. We have manually built a question answering (QA) dataset based on crawled articles and have applied an off-the-shelf QA system based on a fine-tuned BERTimbau Model, achieving an F1-score of 47.0. More importantly, we explored how the proper visualization of attention weights can support helpful interpretations of the system's answers, which is critical in real environments.*

1. Introduction

The Brazilian maritime territory that is often called the "Blue Amazon" has received increasing attention due to its strategic, economic and environmental importance. With an area of 3.6 million km², similar in size to the Brazilian share of the Amazon rainforest, the Blue Amazon borders the entire Brazilian coast and enters the Atlantic Ocean towards Africa [Wiesebron 2013]. From a strategic and economic point of view, the region is home to enormous reserves of oil and minerals [Thompson and Muggah 2015]. On the environmental side, the region holds enormous marine biodiversity. However, despite the growing interest of authorities and companies in the region, the literature on the subject, as well as the knowledge of the population about the region, is still scarce [Castro et al. 2017, Wiesebron 2013]. Hence, it is valuable to develop intelligent tools that have the potential to cover these gaps and improve the knowledge access on the subject.

Amongst the possible initiatives to address this issue, Question Answering (QA) systems offer to a very straightforward strategy: a model is trained in a domain-specific dataset and, once it reaches certain quality criteria, it can serve real users, directly answering their queries in natural language. Some of the best QA systems have relied on the Transformer architecture [Vaswani et al. 2017], often based on BERT [Devlin et al. 2019] or T5 [Raffel et al. 2019].

Despite its success, a Transformer model is a neural model and, as such, it works as a "black box": it is not transparent how the model gets to a predicted output, which may be totally wrong – and this might be potentially harder in today's most popular models with their hundreds of millions or even billions of parameters [Ghaeini et al. 2020]. To meet this challenge, some research has focused on understanding whether BERT learns linguistic features such as part of speech (POS), dependency relationships (DEP) or named entity recognition (NER) on its multiple attention heads [Lample et al. 2016]. Others try to directly visualize the attention weights of models trained in neural machine translation tasks in order to facilitate the users' understanding of the generated translations, especially when they do not make sense [Lee et al. 2017]. Similarly, [Ghaeini et al. 2020] uses the attention weights to interpret decisions from natural language inference (NLI) models, and [Vig 2019a] found even evidence of model biases with the help of them, which is noteworthy for several current practical reasons [Bender and Friedman 2018].

Although there are limitations about the explanatory [Wiegreffe and Pinter 2020] and even the interpretive [Serrano and Smith 2020] power on the importance of attention weights in model decisions, we still argue that intelligent exploration of attention patterns can be instructive and help to understand critical issues or bugs that occur in real systems in production.

To test whether and how a model can adequately answer questions about the Blue Amazon domain, we first manually built a QA dataset based on Wikipedia articles and a textbook [Salas et al. 2011] on the theme with 400 question-answer pairs (QA-pairs).[1] We then applied an off-the-shelf BERTimbau-Base model [Guillou 2021, Souza et al. 2020a] fine-tuned on the Stanford Question Answering Dataset (SQuAD) [Rajpurkar et al. 2016] previously translated into Portuguese. Finally, we used BertViz to explore the views of the Transformer's attention heads, showing how this approach can be useful to understand, parse and debug the model.

2. Background

2.1. Transformers

Transformer have been widely used in several natural language processing applications, such as summarization [Liu and Lapata 2019], question answering [Lewis et al. 2020, Guu et al. 2020] and sentiment analysis [Li et al. 2019].

These neural models are based on the attention mechanism and dispense recurrent neural networks. They are often pre-trained in a self-supervised manner over a huge open-domain database, such as the entire Wikipedia, and then fine-tuned on a specific domain dataset, which can be annotated. The transfer of learning from the first general stage to the second, the specific one, is primarily responsible for these models consistently dominating the rankings of classic natural language processing (NLP) tasks such as the SQuAD dataset [Raffel et al. 2019].

2.2. SQuAD

SQuAD is one of the first large-scale QA datasets with more than one hundred thousand QA pairs. The questions come with a context, and the answer is a segment from it, thus being a reading comprehension dataset. The contexts were taken from Wikipedia articles in the same way as our dataset on the Blue Amazon. As both have the same origin, it is expected that models trained for the SQuAD also present good results in our QA dataset.

[1]The dataset and related codes are available at `https://github.com/C4AI/blab-qa-viz`

3. Models

3.1. BERTimbau fine-tuned on SQuAD in Portuguese

We applied a fine-tuned BERTimbau-Base model [Souza et al. 2020b], a BERT model[2] initially pre-trained in Portuguese, fine-tuned [Guillou 2021] on a translated version of the SQuAD v1.1 dataset. It is currently one of the best available models on the SQuAD dataset in Portuguese.[3]

3.2. BertViz

BertViz was introduced by [Vig 2019a] as an open-source tool for visualizing multi-headed attention in Transformers models, such as BERT and GPT-2. It offers three visualization scales: over the attention-head (Head View), the model (Model View) and the neurons (Neurons View). In this work we focus on the first two ones[4].

It is important to highlight that BertViz offers robust and clear visualizations of the Transformers' attention layers, and represents the state-of-the-art among their peers. The tool intends to support the interpretability of the models, allowing a deeper exploration of how the model arrived at a decision – although a more detailed analysis of such answer is still a step that requires more specialized knowledge.

3.2.1. Head View

Applying the attention-head view in a transformer layer, this tool allows us to visualized the attention patterns produced by the attentions heads. As observed at the Figure 1 [Vig 2019b] there is a example of this view for the following input sentences: at the first image on the left "the cat sat on the mat" and at the image of the center "the cat lay on the rug". The self-attention in this view is represented as lines connecting the tokens that are attending (left) with the tokens being attended to (right). The colors, in this case, are to identify the corresponding attention head(s), and the line weight reflects the attention score. It permits us to select the layer and one or more attention heads represented by the colored patches at the Figure 1, as well filter attention from the combination of sentences (sentence A to sentence A, sentence A to sentence B, sentence B to sentence B, and sentence B to sentence A). In the end is possible to filter attention by token (Figure 2), in that case the target tokens are highlighted by attention strength.

The attention-head view aims to show us how attention browse between tokens for a particular layer and head. In the image on the left of Figure 1, for example, we can see that attention is distributed evenly (not all) across words in the same sentence. In the image on the right of Figure 1, in contrast, we observe that the attention is focused on related words within the opposite sentence.

[2] The main idea behind BERT is the use of pre-training with self-supervised learning, in which the model tries to predict randomly masked words from random sentences taken from the corpus, e.g., a large amount of text gathered in the language in question. This allows the model to store a lot of relevant information of the language being trained within its parameters. It has proven to be very efficient in NLP and it shows state-of-the-art results in several datasets.

[3] The original model can be found here `https://huggingface.co/pierreguillou/bert-base-cased-squad-v1.1-portuguese`

[4] The model are available at `https://github.com/jessevig/bertviz`

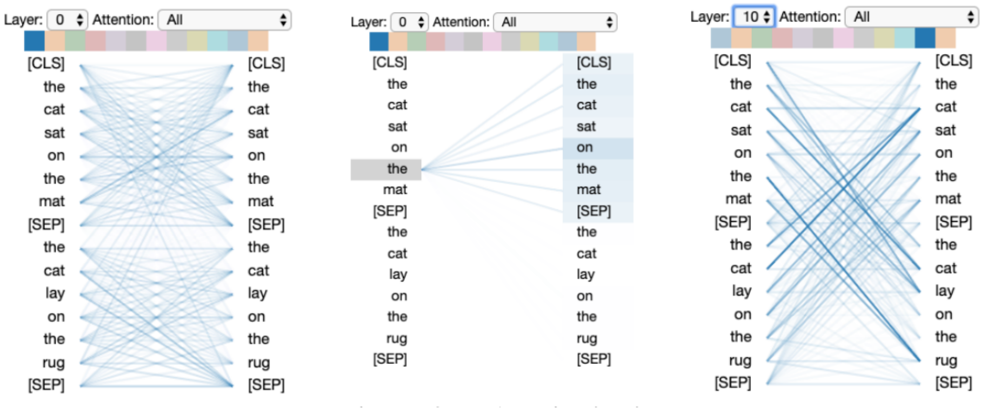

Figure 1. Head view: Attention head view for the sentences "the cat sat on the mat" and "the cat lay on the rug".[Vig 2019b]

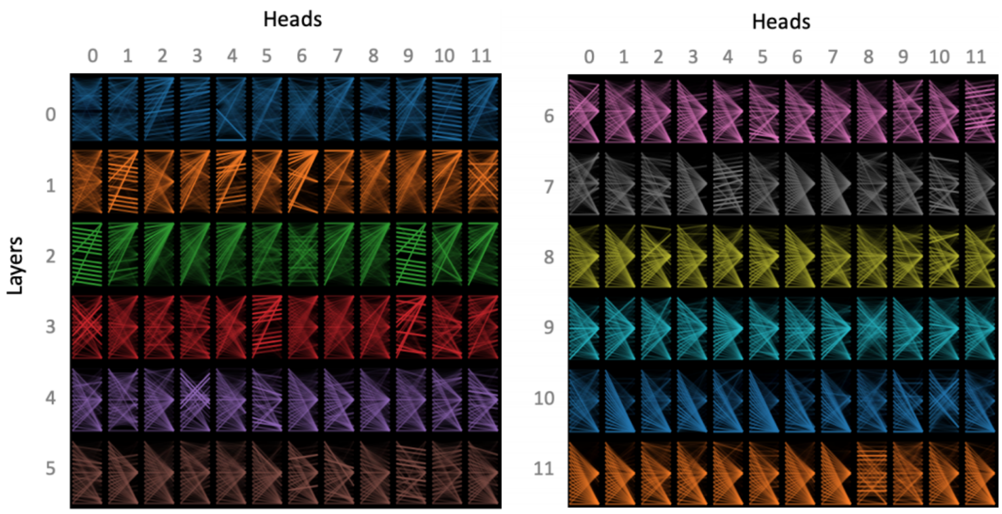

Figure 2. Model view: All attention heads and layers connections are displayed in a single view. While early layers capture low-level patterns, final layers capture higher level sentence representations from separators. [Vig 2019b]

3.2.2. Model View

Providing a view of attention to all models (layers and heads), as we can see in Figure 2 at both images from work of [Vig 2019b]. The author shows the attention heads in tabular form, in each row in the image representing the layers and the columns representing attention heads. As quoted by the author: "Each layer/head is visualized in a thumbnail form that conveys the coarse shape of the attention pattern, following the small multiples design pattern".

This view allows us to browse over all the attention heads and, through the layers of the BERT model, see patterns evolve in the attention heads. In the Figure 2 watching the image on the left, we see that layers 0 to 2 capture low-level patterns, for example, layer 2 and head 0 the attention are to the next word in the sentence. In layers 9 to 11 as we see in the Figure 2 at the image on the right, the attention focus on sentence separators, which that in the author's note it may encode higher level sentence representations [Vig 2019b].

Interpretability of a model is essential for us to detect possible failures or weaknesses in the patterns learning processes, especially in finding dataset characteristics used for learning that can be critical and that possibly will "not converge" to a solution that can be useful by the end-user. Therefore, it is of great value in validating the learning of models and dataset correction.

4. Blue Amazon Dataset Generation

To perform tests on our data about the Blue Amazon, we created the Blue Amazon QA dataset. The Blue Amazon QA dataset is fully composed of handcrafted QAs about the context extracted from Wikipedia articles related to the Blue Amazon.

We applied the following methodology to create the dataset: Firstly, we defined the Wikipedia articles that would be utilized as context. We gathered 40 articles related to the Brazilian coast geography, oceanic islands, marine biodiversity, coastal cities and villages, among others. Then, we created a Telegram bot assistant that interacted with users upon solicited, displaying random paragraphs from the set of Wikipedia articles and asked for QAs about the topics. The bot would save the QAs sent by the users, automatically associating them to the corresponding Wikipedia article and paragraph position.

We assigned four specialists to generate QAs, that were advised to generate complete and self-explanatory questions, accompanied by direct short answers. The specialists created 100 QA pairs each, following the same methodology. After the generation cycle, all resulting QAs were reviewed by the authors, and the resulting dataset had around 400 QAs about the Wikipedia extracts.

Table 1 shows some examples of QAs with their associated Wikipedia article title (Wiki. ref.) and reference paragraph position (Pos.). As one can see, most of the dataset is composed of direct questions accompanied by a maximum of three words per answer.

Table 1. Extract of the handcrafted Blue Amazon QA dataset. The dataset has question and answers associated to a specific Wikipedia article (Wiki. ref.) and paragraph position (Pos.).

Wiki. ref.	Question	Answer	Pos.
Coastline of Brazil	Qual é o tamanho da costa do Amapá? (*What is the length of the Amapá's coastline?*)	Quase 600 km (*Almost 600 km*)	14
Coastline of Brazil	Quantas praias tem a costa do Amapá? (*How many beaches are in Amapá's coastline ?*)	3 praias (*3 beaches*)	12
Port of Tubarão	Quando o Porto de Tubarão foi criado? (*When was the Port of Tubarão created?*)	1966	0
Rocas Atoll	Qual é o único atol do atlântico sul? (*What is the only South Atlantic atoll?*)	Atol das Rocas (*Rocas Atoll*)	0
Guanabara Bay	Qual é o tamanho da baía de Guanabara? (*How long is guanabara bay?*)	31 km	1

5. Results

5.1. Model Performance on the Blue Amazon Dataset

In table 2, we report the F1-score, the Exact Match (EM), and the RougeL (R-L) metrics obtained with our fine-tuned BERTimbau model on the Blue Amazon Dataset.

Table 2. Result of QA dataset test with the metrics.

Model	F1-Score	EM (Exact Match)	RougeL (R-L)
BERT-Squad	47.0	10.0	48.4

To perform the experiments, we leveraged a machine in Colab Research Google with GPUs Tesla K80 12GB, 12.69GB of RAM and 73.27GB of Disk. Under these conditions, the inference time – that is, the time it took to a model to generate an answer to a single question from the test set – was around 0.52s.

5.2. Case Study: Interpretability on Blue Amazon Portuguese QA

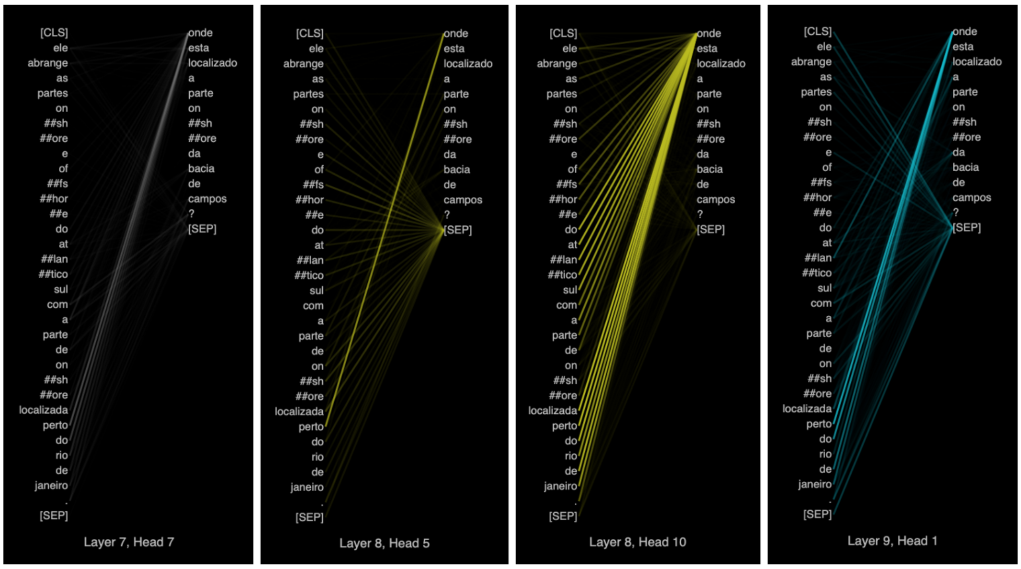

Figure 3. Model view: Attention head view for the sentences *Ele abrange as partes onshore e offshore do Atlantico Sul com a parte de onshore localizada perto do Rio de Janeiro* (*"It covers the onshore and offshore parts of the South Atlantic with the onshore part located near Rio de Janeiro"*).

It is easy to be amazed or utterly disappointed with the result of a complex model like BERT for QA or many other Transformers models. On one hand, we are intrigued by the surprising result of getting the correct answer for one's question; on the other hand, we wonder what is wrong in the model when it does not get the expected result and we are also stuck because it is not easy to control the learning of these models to the point of optimizing until it successfully answers all the questions in a domain of knowledge (Utopia). Worse than that is not being able to explain an answer the result to a researcher or user. So, interpretability methods are critical because they give us the ability to interpret a case as a QA model based on BERT and Squad with knowledge about our Blue Amazon that we are interested in improving here.

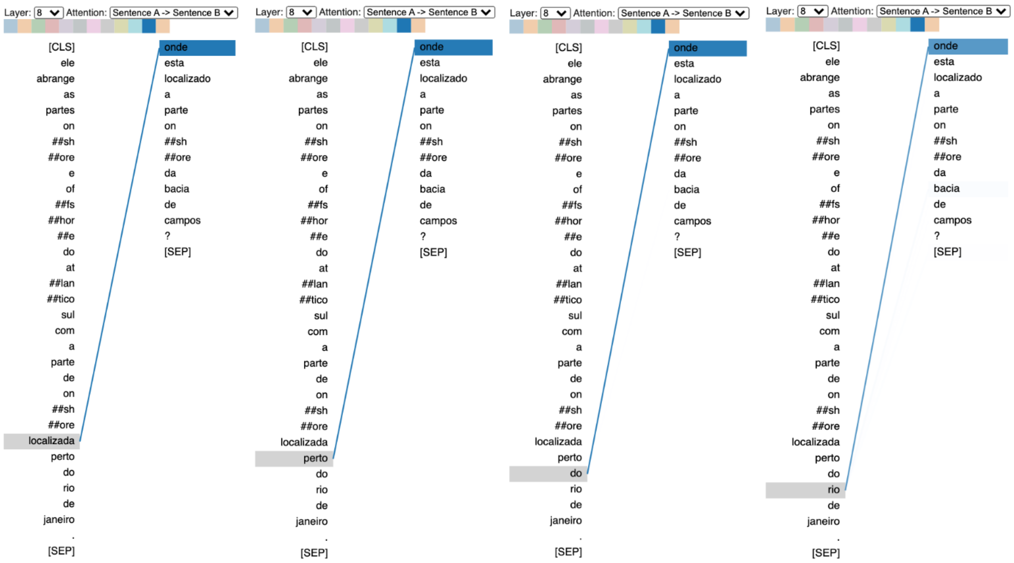

Figure 4. Head view: Attention head view for the sentences *Ele abrange as partes onshore e offshore do Atlantico Sul com a parte de onshore localizada perto do Rio de Janeiro* (*"It covers the onshore and offshore parts of the South Atlantic with the onshore part located near Rio de Janeiro"*).

Figure 5. Head view: Attention head view for the sentences *Ele abrange as partes onshore e offshore do Atlantico Sul com a parte de onshore localizada perto do Rio de Janeiro* (*"It covers the onshore and offshore parts of the South Atlantic with the onshore part located near Rio de Janeiro"*).

In our scenario, in addition to an answer given by the algorithm or model, a user also wants answers to questions such as: Why did the model give that answer? How did it learn the language of that text? or a simply Why?. In this context, to clarify so many possible questions raised by users, as a case of study a question was asked to the QA model about Blue Amazon like: *Onde está localizado a parte onshore da bacia de Campos?* (*"Where is the onshore part of the Campos basin located?"*). The text snippet

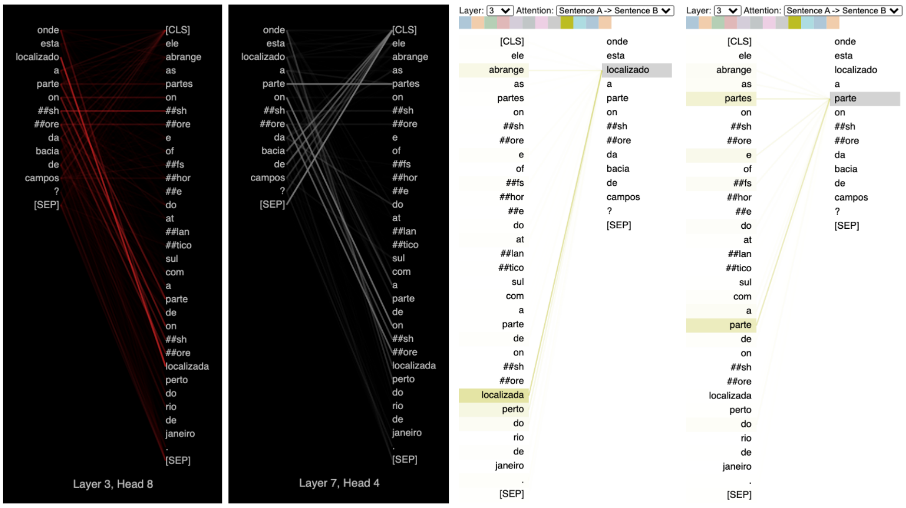

Figure 6. Head and Model view: Attention head view for the sentences *Onde está localizado a parte onshore da bacia de Campos?* (*"Where is the onshore part of the Campos basin located?"*).

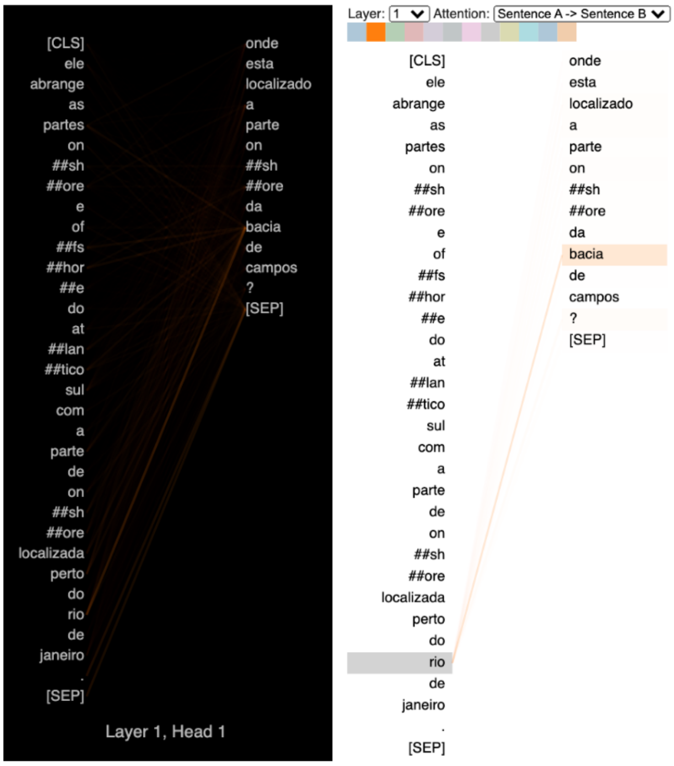

Figure 7. Head and Model view: Attention head view for the sentences *Onde está localizado a parte onshore da bacia de Campos?* (*"Where is the onshore part of the Campos basin located?"*).

containing in the knowledge to this question can be observed here: *Ele abrange as partes onshore e offshore do Atlantico Sul com a parte de onshore localizada perto do Rio de Janeiro* (*"It covers the onshore and offshore parts of the South Atlantic with the onshore part located near Rio de Janeiro"*). Finally, the answer generated was for the question is:

perto do Rio de Janeiro (*“near Rio de Janeiro”*).

In order to interpret the output of the model, the BertViz library was used to look at the relationships between the sentences in pieces analyzing each layer and heads of attention. Bringing up the view of the heads of attention in Figure 3, Sentence A (left) → Sentence B (right). It is possible to see the attention given to the word *onde* (*“where”*) connected to the words *perto do Rio de Janeiro* (*“near do Rio de Janeiro”*, indicating a relationship between the answer and the question, besides suggesting that the model in a given layer is able to understand semantic relationships between words as we observed between the adverbs *onde* (*“where”*) and *perto* (*“close”*) or between adverb and verb *onde* (*“where”*) and *localizada* (*“located”*). It is easy to notice that as you go to the next layer, this relationship becomes more and more intense, and the attention is stronger. In Figure 4 and Figure 5 we show this strong connection by visualizing word by word.

Another common pattern observed in this model is the layers paying attention to the same words, same roots or close semantic meanings, as is the case in Figure 6, in which we see attention to the word *localizada* (*“located”*) with the word *localizado* (*“located“*), the word *abrange* (*“embrace”*) with the word *localizado* (*“located”*), the word *partes* (*“parts”*) with *parte* (*“part”*) and finally the word *onshore*.

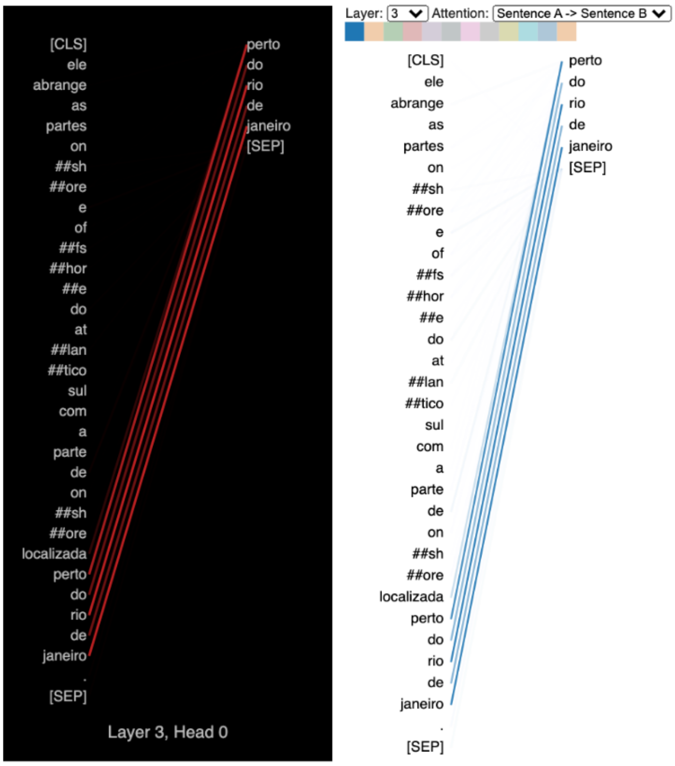

Figure 8. Head and Model view: Attention head view for the sentences *perto do Rio de Janeiro* (*“near Rio de Janeiro”*).

Although *Rio de Janeiro* is the name of a state in Brazil, it is important to highlight that the sentences undergo a pre-treatment before they are fit to the model, in which case all the words are in their uppercase form. This lends a different meaning to the words *rio* (*“river”*) and *janeiro* (*“january”*), so the relation of attention observed between the word *rio* (*“river”*) and *bacia* (*“basin”*) in Figure 7 makes perfect sense. It also shows how attention based models are powerful for language understanding, unlike word frequency based models like BM25 [Robertson and Zaragoza 2009] that only get relationships if the words are identical.

Finally, we note the attention given to the words that together are the answers to the question asked. We can see in Figure 8 that all attention was given to the same words in both sentences, with virtually no attention given to another word in the sentence containing the answer, which strongly indicates the correct classification of the models when deciding that the correct answer to the question *Onde está localizado a parte onshore da bacia de Campos?* (*"Where is the onshore part of the Campos basin located?"*) it is, in fact, *perto do Rio de Janeiro* (*"near Rio de Janeiro"*).

6. Conclusion

In this work, we combined an off-the-shelf BERTimbau model fine-tuned on a version of the SQuAD dataset translated into Portuguese with BertViz to visualize attention weights. It proved to be a useful way to complement the interpretation of the model's decisions on the generated answers. Also, we contributed with a QA dataset built from filtered articles from Wikipedia and a textbook, both in Portuguese, about the Blue Amazon. As future work, we would like to expand the dataset on the Blue Amazon by creating new questions manually and by filtering and translating domain-specific QA-pairs from massive open-domain datasets, such as PAQ [Lewis et al. 2021] or Natural Questions [Kwiatkowski et al. 2019]. In addition, we will focus our efforts on developing an interpretability method that integrates the two BertViz's views and exposes the value vectors beyond the assertions, giving us a straightforward conclusion about the results in a user-friendly layout.

Acknowledgements

This work was financed in part by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES, Finance Code 001) and by the *Itaú Unibanco S.A.*, through the *Programa de Bolsas Itaú* (PBI) of the *Centro de Ciência de Dados* (C^2D) of *Escola Politécnica* of *Universidade de São Paulo* (USP). We also gratefully acknowledge support from *Conselho Nacional de Desenvolvimento Científico e Tecnológico* (CNPq) (grants 312180/2018-7 and 310085/2020-9) and the *Center for Artificial Intelligence* (C4AI-USP), with support by the *Fundação de Amparo à Pesquisa do Estado de São Paulo* (FAPESP, grant 2019/ 07665-4) and by the *IBM Corporation*. Any opinions, findings, and conclusions expressed in this manuscript are those of the authors and do not necessarily reflect the views, official policy or position of the Itaú-Unibanco, IBM Corporation, CAPES, FAPESP, and CNPq.

References

Bender, E. M. and Friedman, B. (2018). Data Statements for Natural Language Processing: Toward Mitigating System Bias and Enabling Better Science. *Transactions of the Association for Computational Linguistics*, 6:587–604.

Castro, B. M., Brandini, F. P., Dottori, M., and Fortes, J. F. (2017). A Amazônia Azul: recursos e preservação. *Revista USP*, (113):7.

Devlin, J., Chang, M. W., Lee, K., and Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1(Mlm):4171–4186.

Ghaeini, R., Fern, X. Z., and Tadepalli, P. (2020). Interpreting recurrent and attention-based neural models: A case study on natural language inference. *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018*, pages 4952–4957.

Guillou, P. (2021). Portuguese bert base cased qa (question answering), finetuned on squad v1.1.

Guu, K., Lee, K., Tung, Z., Pasupat, P., and Chang, M. (2020). REALM: retrieval-augmented language model pre-training. *CoRR*, abs/2002.08909.

Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Kelcey, M., Devlin, J., Lee, K., Toutanova, K. N., Jones, L., Chang, M.-W., Dai, A., Uszkoreit, J., Le, Q., and Petrov, S. (2019). Natural questions: a benchmark for question answering research. *Transactions of the Association of Computational Linguistics*.

Lample, G., Ballesteros, M., Subramanian, S., Kawakami, K., and Dyer, C. (2016). Neural architectures for named entity recognition. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.

Lee, J., Shin, J. H., and Kim, J. S. (2017). Interactive visualization and manipulation of attention-based neural machine translation. *EMNLP 2017 - Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Proceedings*, pages 121–126.

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., and Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive nlp tasks. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M. F., and Lin, H., editors, *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.

Lewis, P., Wu, Y., Liu, L., Minervini, P., Küttler, H., Piktus, A., Stenetorp, P., and Riedel, S. (2021). Paq: 65 million probably-asked questions and what you can do with them.

Li, X., Bing, L., Zhang, W., and Lam, W. (2019). Exploiting BERT for end-to-end aspect-based sentiment analysis. *CoRR*, abs/1910.00883.

Liu, Y. and Lapata, M. (2019). Text summarization with pretrained encoders. *CoRR*, abs/1908.08345.

Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Peter, W. L., and Liu, J. (2019). Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv*, 21:1–67.

Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. (2016). Squad: 100, 000+ questions for machine comprehension of text. *CoRR*, abs/1606.05250.

Robertson, S. and Zaragoza, H. (2009). *The probabilistic relevance framework: BM25 and beyond*, volume 3.

Salas, S., Chuenpagdee, R., Charles, A. T., Seijo, J. C., et al. (2011). *Coastal fisheries of Latin America and the Caribbean*, volume 544. Food and Agriculture Organization of the United Nations ˆ eRome Rome.

Serrano, S. and Smith, N. A. (2020). Is attention interpretable? *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of the Conference*, pages 2931–2951.

Souza, F., Nogueira, R., and Lotufo, R. (2020a). BERTimbau: Pretrained BERT Models for Brazilian Portuguese. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 12319 LNAI, pages 403–417.

Souza, F., Nogueira, R., and Lotufo, R. (2020b). BERTimbau: Pretrained BERT Models for Brazilian Portuguese. In Cerri, R. and Prati, R. C., editors, *Intelligent Systems*, pages 403–417, Cham. Springer International Publishing.

Thompson, B. N. and Muggah, R. (2015). The Blue Amazon: Brazil Asserts Its Influence Across the Atlantic.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 2017-Decem(Nips):5999–6009.

Vig, J. (2019a). A multiscale visualization of attention in the transformer model. *ACL 2019 - 57th Annual Meeting of the Association for Computational Linguistics, Proceedings of System Demonstrations*, pages 37–42.

Vig, J. (2019b). Bertviz: A tool for visualizing multi-head self-attention in the bert model.

Wiegreffe, S. and Pinter, Y. (2020). Attention is not not explanation. *EMNLP-IJCNLP 2019 - 2019 Conference on Empirical Methods in Natural Language Processing and 9th International Joint Conference on Natural Language Processing, Proceedings of the Conference*, pages 11–20.

Wiesebron, M. (2013). Amazônia Azul: Pensando a Defesa Do Território Marítimo Brasileiro. *AUSTRAL: Brazilian Journal of Strategy & International Relations*, 2(3):107–132.