

Multilingual, Multi-scale and Multi-layer Visualization of Sequence-based Intermediate Representations

Carlos Escolano^{*◇}, Marta R. Costa-jussà^{*◇}, Elora Lacroux[◇] and Pere-Pau Vázquez^{*◇}

^{*} TALP Research Center, [◇]Universitat Politècnica de Catalunya, Barcelona

^{*} ViRVIG Group

{carlos.escolano, marta.ruiz}@upc.edu

lacrouxelora@gmail.com, pere.pau@cs.upc.edu

Abstract

The main alternatives nowadays to deal with sequences are Recurrent Neural Networks (RNN), Convolutional Neural Networks (CNN) architectures and the Transformer. In this context, RNN's, CNN's and Transformer have most commonly been used as an encoder-decoder architecture with multiple layers in each module. Far beyond this, these architectures are the basis for the contextual word embeddings which are revolutionizing most natural language downstream applications.

However, intermediate layer representations in sequence-based architectures can be difficult to interpret. To make each layer representation within these architectures more accessible and meaningful, we introduce a web-based tool that visualizes them both at the sentence and token level. We present three use cases. The first analyses gender issues in contextual word embeddings. The second and third are showing multilingual intermediate representations for sentences and tokens and the evolution of these intermediate representations along the multiple layers of the decoder and in the context of multilingual machine translation.

1 Introduction

The Transformer (Vaswani et al., 2017) is a powerful architecture that was initially proposed to train neural machine translation. This architecture deals with variable sequences by concatenating feed-forward networks and attention-based mechanisms. While the composed modules of the Transformer may not be complex by themselves, it is the composition of several layers of these modules that make the architecture less interpretable.

We are aiming at providing a tool to give insights to the sentences and token representation from each layer in the Transformer. Far beyond the Transformer interpretation which has become

by de-facto the state-of-the-art in machine translation, our tool is able to represent intermediate representations of other sequence-based architectures such as RNNs (Bahdanau et al., 2014) or ConvS2S (Gehring et al., 2017) as well. Note that sequence-based architectures are having impact in many multimodal applications such as image captioning and speech recognition (Kaiser et al., 2017; Chan et al., 2016).

The uses of our visualization tool are quite a few varying from social bias, multilingual or linguistic analysis. In particular, we focus in analysing the gender inequalities in contextual word embeddings and the common language representation in a multilingual machine translation system.

2 Visualization tool

In this section we present a multi-scale and multi-layer visualization tool for the sequence-based architectures, available as tool¹ and as a demo². The tool is implemented in Python using the Bokeh library for data visualization and the Flask library as web microframework to embed the Bokeh dashboards on the webpage.

The tool consists in using as input fixed-representations, being a matrix of dimensions the embedding size per sentence length (in tokens). Therefore, the input data required are the sentences to be represented (*txt*), the sentence representations (*json*) and optionally the tokens embeddings (*json*). Then, a UMAP (McInnes et al., 2018) dimensionality reduction is performed to plot the representation of this multidimensional data in two dimensions. This dimensionality reduction is performed for the fixed-representations at the sentence and token level. The tool comprises two views: multi-scale intermediate repre-

¹<https://github.com/elorala/interlingua-visualization>

²<https://upc-nmt-vis.herokuapp.com/>

sensation for one layer and multi-layer sentence representation. These three views can be either monolingual or multilingual. The main page of the tool comprises these three views for the user to choose.

We describe these three views on different use cases. For the first view, we show the use cases of detection of gender bias in contextual word embeddings and common representation in multilingual machine translation. For the second view, the use case builds on layer interpretation of multi-way parallel sentences in a translation decoder and showing which layer carries out higher semantic meaning.

2.1 Multi-scale Intermediate representation

This visualization consists of two coordinated views, that encode different information through scatterplots. The one on the left shows the M sentence intermediate representations. Each dot in the sentence graph corresponds to one sentence, by hovering on a point we visualize the sentence as well as the arrows to the corresponding translation sentences, in case we are working with multilingual data. There is an option to visualize a particular sentence by writing it in the search bar. The search bar has an autocomplete feature (activated when typing two characters) and then, the user can click on the right suggestion.

The right view shows the tokens. Initially, when no sentence from the previous view is selected, this plot shows all vocabulary tokens. By brushing over one or more sentences (in left view), the right view filters out the tokens *not belonging* to the selected sentence (and the tokens that compose the parallel sentences in the other languages). Once the user selects a sentence by clicking or searching, only the words from this sentence (and its translations) remain on the chart. By hovering on a point, the user can see the text of the word, analogously to the sentences view.

Sentences and tokens can be simultaneously visualized for all languages that we are studying and we can interpret the intermediate representation in terms of both granularity levels. See Figures 4 and 5 which are as well examples of the second use case (explained as follows).

Use case 1: Gender bias in Contextual Word Embeddings. The objective of this use case is to visualize the contextual word representations on a set of occupational vocabulary. We use the

ELMO implementation (Peters et al., 2018), based on RNNs and as data, we use 1019 sentences from previous work (Font and Costa-jussà, 2019) that follow the next template *I've known him/her for a long time, my friend works as a occupation*. Examples of occupations include: *accounting clerk, nurse midwife or biological scientist*. Since we have two sets: one for female templates and another for male templates, we use the two sets as if they were different languages. We visualize 2-dimensional representations of sentences and words. For sentences (see Figure 1), we see that sentences with similar professions (i.e. *financial manager, personal financial advisor*) tend to be close in the space for both female and male versions. However, when visualizing words, in the case of *financial manager*, words for female and male representation are placed in very distant points in the space as seen in Figure 2. On the contrary, words for female and male representation in the case of *personal financial advisor* are represented together as seen in Figure 3. So, we conclude that *financial* in a male/female context is differently represented if attached to *manager* but the same *financial* is similarly represented in male/female context if attached to *personal* and *advisor*. Our tool allows to visualize that contextual word embeddings encode gender biases and this conclusion is coherent with previous literature experiments (Basta et al., 2019).




Figure 1: Contextual word embedding representation at the sentence level (sentences *I've known him/her for a long time, my friend is a financial manager/personal financial advisor*). Sentences referring to males are in green, sentences referring to females are in red.

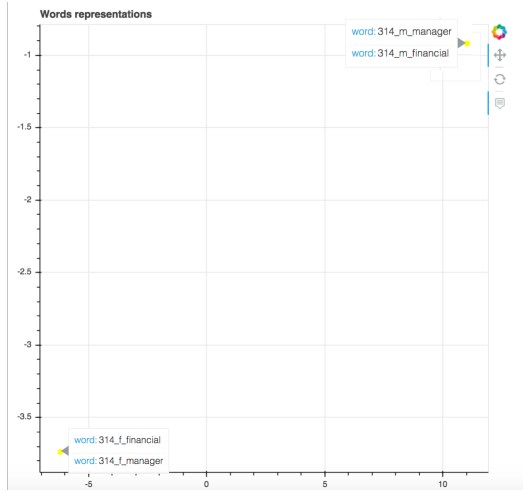


Figure 2: Contextual word embedding representation at the token level *financial manager*.




Figure 3: Contextual word embedding representation at the token level *personal financial advisor*.

Use case 2: Multilingual common representation in translation. Nowadays, there are two main architectures for multilingual neural machine translation which are a universal shared encoder and decoder and independent multiple encoders and decoders. In both cases, there is an intermediate representation where sentences that have similar meanings should be represented close in the space. For our second and third use case, we use the intermediate representations of the multilingual Transformer-based architecture presented in (Escolano et al., 2019). Basically, the architecture consists in independent encoders and decoders with a forced-interlingua space. This system is trained on data extracted from the UN (Ziemski et al., 2016) and EPPS datasets (Koehn, 2005) that provide 15 million parallel sentences

between English and Spanish and French. *newstest2012* and *newstest2013* were used as validation and test sets, respectively. These sets provide parallel data between the 3 languages.

Figure 4 shows 130 sentences extracted from the test set, in the 3 languages at hand and in the common space (at the output of the encoder). When we select a particular sentence (e.g. *people accept orders*.), for each token in the sentence selected, the user can select to visualize the token representations (e.g. *people*) as shown in Figure 5. From this visualization we conclude that the model is not able to group together sentences with the same meaning across languages.

2.2 Multi-layer sentence representation

This visualization shows T layers simultaneously for single or multiple languages in a small multiples design. This facilitates the analysis of sentence representation evolution across all the layers of the Transformer at once. See Figure 6.

On each view, we can display the sentence by hovering. In order to emphasize the distances between the translations and to have a better insight of the evolution, the link between the most dissimilar are displayed on the plots. By hovering on the lines, the user can obtain the cosine distance value computed on SciPy. On the views, only the distances superior to 1 are displayed. Even if the dimensionality reduction of UMAP does show interpretable distances (McInnes et al., 2018), showing consecutive layers of the Transformer, and seeing the evolution of the representations allows us to draw hints about the layer roles as we will see in the third use case.

Finally, the tool allows for analysis in multiple layers and languages. This means that initially, the multiple layers represented on the dashboard are in one particular language. However, the user can switch to the multiple layers from another language by using the selection tool at the top of the page. Since all views are synchronized, upon changing the language set, all of them change accordingly.

Use case 3: Multilingual Layer Interpretation in Translation Decoding Encoders and decoders in a neural machine translation system are usually composed of different layers. The role of each layer is difficult to interpret. Visualizing sentences at each of these layers can help us on identifying the sentence distance evolution giving




Figure 4: Multilingual common representation at the sentence level (sentence *people accept orders .*). English in red, Spanish in green and French in blue.




Figure 5: Multilingual common representation at the token level (token *people*).

us hints of different linguistic roles for the layers when compared between them.

In the current example, we are representing the same set and architecture as in use case 2, but for the 6 decoder layers. Figure 6 shows the plot for these layers and Figure 7 shows how it performs hovering on a point (e.g. showing sentences, *unexpected consequences.*, right) and hovering on a line (e.g. showing distance measure, left). Since we show sentences with the same meaning in different languages, we interpret that the layer that tends to better cluster sentences compared to contiguous layers is the one with higher semantic implications. From Figure 6, we conclude that higher layers in the decoder (specially 4 and 5) better group sentences (see axes values).

3 Adaptability

In this paper, we have discussed three use cases. However, our tool is highly flexible and adaptable,

and it allows for a large variety of tasks. The system only requires data to be formatted as a *JSON* file following the structures defined in Figure 8.

The structure from use cases 1 and 2 defines the relation between sentence and token representations. For each token and embedding a 2-dimensional is defined, showing its coordinates in the final plots.

On the other side, the structure from use case 3 contains the representations of the layers to be plotted and it is described as an array containing the coordinates for each sentence.

This implementation allows our tool to be agnostic to factors such as vocabulary sizes and dimensionality reductions techniques, as they are applied before *JSON* creation.

4 Related Work

Given the versatility of the sequence architectures, the current tool feeds from vast research areas including contextual word embeddings, multilingual models, visualization and interpretability of sequence models, zero-shot learning. However, we just refer here to the closest and recent works.

Gender bias. Gender bias has recently been analysed in contextual word embeddings (Zhao et al., 2019; Basta et al., 2019). Our tool aims at following-up this kind of research to work towards techniques that are able to neutralize these and other social biases.

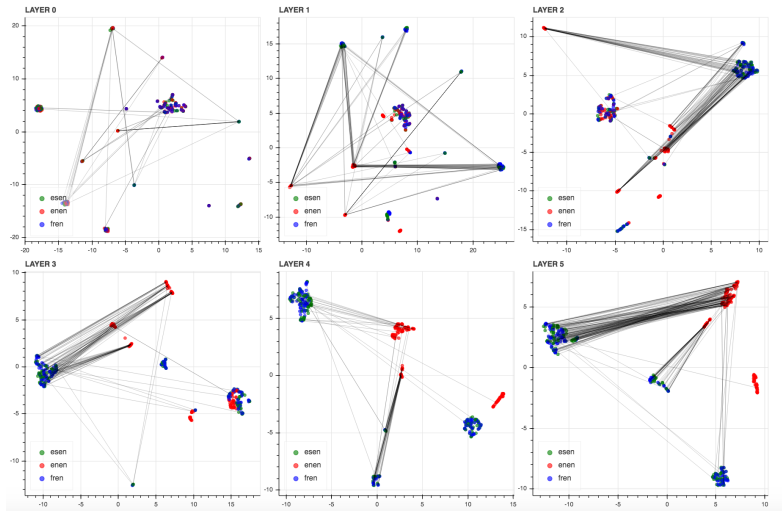


Figure 6: Decoder layer multilingual sentence representation.




Figure 7: Decoder layer multilingual sentence representation: distance and sentences.

```

{
  "content": {
    "id": {
      "sentence": "",
      "embedding": "",
      "words": {
        "word1": "",
        "wordN": ""
      }
    }
  }
}

{
  "type": "UMAP",
  "content": [
    {
      "input_source0": "",
      "input_sourceN": ""
    },
    {
      "input_source0": "",
      "input_sourceN": ""
    }
  ]
}

```

Figure 8: Required JSON structures: (Left) use cases 1 and 2 and (Right) use case 3

Multilinguality analysis. It is quite a common practice to visualize intermediate representations of sequence-to-sequence models (Johnson et al., 2017; Escolano et al., 2019). Our tool is not limited to this sentence representation of the intermediate representation, but it also includes the token-level representation. By simultaneously providing this two-granularity level representation we are aiming at a deeper analysis for monolingual,

cross-lingual and multilingual natural language processing downstream applications in general.

Linguistic insights. (Raganato and Tiedemann, 2018) show interesting findings about syntactic and semantic behavior across Transformer layers. Following this research line, our tool can further analyse how similar sentences in multiple languages evolve in their intermediate layer representations as well as monolingual sentences with same syntactic or morphological patterns.

Finally, regarding related visualizations and demonstrations, authors in (Li et al., 2016) make an visual analysis of neural models specifically in natural language processing (but focusing on previous architectures to the Transformer), while (Vig, 2019) analyse the attention in the Transformer at multiple-scales and show different use cases on contextual word embeddings. Our tool further adds to these previous works by focusing on the intermediate representations.

5 Conclusions

We have presented an extremely flexible and adaptable visualization tool for multilingual intermediate representations of text both at the sentence and token’s level. Together with our tool we have presented three use cases in the context of gender bias analysis in contextual word embeddings and for multilingual intermediate representations of machine translation.

Acknowledgements

Authors want to thank Christine Raouf Basta for sharing her expertise in contextual word embeddings. This work is supported by a Google Faculty Research Award. This work is also supported by the Spanish Ministerio de Economía y Competitividad, the European Regional Development Fund and the Agencia Estatal de Investigación, through the postdoctoral senior grant Ramón y Cajal, contracts TEC2015-69266-P and TIN2017-88515-C2-1-R(GEN3DLIVE) (MINECO/FEDER,EU), and contract PCIN-2017-079 (AEI/MINECO).

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Christine Basta, Marta R. Costa-jussà, and Noe Casas. 2019. Evaluating the underlying gender bias in contextualized word embeddings. In *Proc. of the 1st ACL Workshop on Gender Bias for Natural Language Processing*.
- William Chan, Navdeep Jaitly, Quoc V. Le, and Oriol Vinyals. 2016. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *ICASSP*.
- Carlos Escolano, Marta R. Costa-jussà, and Jos A. R. Fonollosa. 2019. From bilingual to multilingual neural machine translation by incremental training. In *Proc. of the ACL Student Research Workshop*.
- Joel Escudé Font and Marta R. Costa-jussà. 2019. Equalizing gender biases in neural machine translation with word embeddings techniques. In *Proc. of the 1st ACL Workshop on Gender Bias for Natural Language Processing*.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. 2017. Convolutional sequence to sequence learning. In *Proceedings of the 34th ICML - Volume 70*, pages 1243–1252. JMLR.org.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. Google’s multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Lukasz Kaiser, Aidan N Gomez, Noam Shazeer, Ashish Vaswani, Niki Parmar, Llion Jones, and Jakob Uszkoreit. 2017. One model to learn them all. *arXiv preprint arXiv:1706.05137*.
- Philipp Koehn. 2005. Europarl: A parallel corpus for statistical machine translation. In *MT summit*, volume 5, pages 79–86.
- Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2016. 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.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Grossberger. 2018. Umap: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3(29):861.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the NAACL*, pages 2227–2237, New Orleans.
- Alessandro Raganato and Jörg Tiedemann. 2018. An analysis of encoder representations in transformer-based machine translation. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 287–297, Brussels, Belgium.
- 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. 2019. A multiscale visualization of attention in the transformer model. In *Proc. of the ACL System Demonstrations*.
- Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. Gender bias in contextualized word embeddings. In *Proceedings of the 2019 Conference of the NAACL*, pages 629–634, Minneapolis, Minnesota.
- Michał Ziemski, Marcin Junczys-Dowmunt, and Bruno Pouliquen. 2016. The united nations parallel corpus v1. 0. In *Lrec*.