Erratum to *Encoders Help You Disambiguate Word Senses in Neural Machine Translation*

Gongbo Tang[1] **Rico Sennrich**[2,3] **Joakim Nivre**[1]

[1]Department of Linguistics and Philology, Uppsala University
[2]Institute of Computational Linguistics, University of Zurich
[3]School of Informatics, University of Edinburgh
`firstname.lastname@{lingfil.uu.se, ed.ac.uk}`

Abstract

Tang et al. (2019) explored the ability of NMT encoders and decoders to disambiguate word senses by evaluating hidden states. The work contained an error when evaluating decoder hidden states due to an implementation bug in generating decoder hidden states, which was discovered after the publication. After correcting the error, decoder hidden states from both Transformer and *RNNS2S* models achieve higher accuracy than encoder hidden states, which accords with the hypothesis that decoders provide further relevant information for disambiguation. The error does not otherwise affect the originally reported conclusions.

1 Error

Tang et al. (2019) explored the ability of NMT encoders and decoders to disambiguate word senses by evaluating hidden states. The work contained an error when evaluating decoder hidden states due to an implementation bug in generating decoder hidden states, which was discovered after the publication.

As a result of this error, the accuracy of decoder hidden states (*DEC*) reported in Table 2 does not reveal the ability of decoders in disambiguation. We have re-run the experiments with fixed decoder hidden states. A list of corrections are given below.

2 List of Corrections

- Abstract, the penultimate sentence:

  > In contrast to encoders, the effect of decoder is different in models with different architectures.

  should become

  > Decoders could provide further relevant information for disambiguation.

- Section 1, the third finding

  > Decoders hidden states have different effects on WSD in *Transformer* and *RNNS2S*.

  should be removed.

- Table 2:

	DE→EN		DE→FR	
	RNN.	*Trans.*	*RNN.*	*Trans.*
Embedding	63.1	63.2	68.7	68.9
ENC	94.2	97.2	91.7	95.6
DEC	97.9	91.2	95.1	91.6

  should become

	DE→EN		DE→FR	
	RNN.	*Trans.*	*RNN.*	*Trans.*
Embedding	63.1	63.2	68.7	68.9
ENC	94.2	97.2	91.7	95.6
DEC	97.5	98.3	95.1	96.9

- In Section 3.1, the last sentence:

  > In addition, *DEC* achieves even higher accuracy than *ENC* in *RNNS2S* models but not in *Transformer* models.

  should be

  > In addition, *DEC* achieves even higher accuracy than *ENC* in both *RNNS2S* models and *Transformer* models.

- Section 4.1.3:

  > As Table 2 shows, RNN decoder hidden states could further improve the classification accuracy which accords with our hypothesis. It implies that the relevant information for WSD in the target-side has

> been well incorporated into the decoder hidden states to predict the translations of ambiguous nouns. It is curious that Transformer decoder hidden states are inferior to Transformer encoder hidden states in our WSD classification task, given that Tang et al. (2018) and Rios et al. (2018) report better results with contrastive evaluation and semi-automatic evaluation of 1-best translations for Transformer models than for *RNNS2S*. However, note that our evaluation merely tests whether the information necessary for word sense disambiguation is encoded in hidden states and can be extracted by our binary classifier. In practice, decoder hidden states are used for predicting a target word from the entire vocabulary, and thus need to encode additional information which may confound our classifier.
>
> Despite these differences between *RNNS2S* and the Transformer, our results show that WSD is already possible on the basis of the encoder representation of the ambiguous noun, and that extracting contextual information via encoder-decoder attention or from the target history is not essential for WSD.

should be

> As Table 2 shows, decoder hidden states could further improve the classification accuracy which accords with our hypothesis. It implies that the relevant information for WSD in the target-side has been well incorporated into the decoder hidden states to predict the translations of ambiguous nouns.
>
> Although decoder hidden states achieve higher accuracy than encoder hidden states, the improvement is not as big as that achieved by encoder hidden states over word embeddings. This indicates that most of the disambiguation work is done by encoders.

- In Section 5, the last sentence of the first paragraph:

  > Moreover, the effect of decoder hidden states on WSD is different in Transformer and *RNNS2S* models.

  should be

  > Even though decoders could provide more relevant information for disambiguation, most of the disambiguation work is done by encoders.

- Footnote 6 in the Appendix:

  > The classifiers fed decoder states are trained 200 epochs to converge.

  should be removed.

3 Conclusion

There was an implementation bug in generating decoder hidden states in Tang et al. (2019) and it was discovered subsequent to publication. After correcting the error, decoder hidden states from both Transformer and *RNNS2S* models achieve higher accuracy than those encoder hidden states. The error does not affect the originally reported statement that the encoder is the primary component for disambiguation.

References

Annette Rios, Mathias Müller, and Rico Sennrich. 2018. The word sense disambiguation test suite at WMT18. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 588–596, Belgium, Brussels. Association for Computational Linguistics.

Gongbo Tang, Mathias Müller, Annette Rios, and Rico Sennrich. 2018. Why self-attention? a targeted evaluation of neural machine translation architectures. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4263–4272, Brussels, Belgium. Association for Computational Linguistics.

Gongbo Tang, Rico Sennrich, and Joakim Nivre. 2019. Encoders help you disambiguate word senses in neural machine translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1429–1435, Hong Kong, China. Association for Computational Linguistics.