

Modeling Five Sentence Quality Representations by Finding Latent Spaces Produced with Deep Long Short-Memory Models

Pablo Rivas

Department of Computer Science
School of Computer Science and Mathematics
Marist College, Poughkeepsie, New York, USA
Pablo.Rivas@Marist.edu

Abstract

We present a study in which we train neural models that approximate rules that assess the quality of English sentences. We modeled five rules using deep LSTMs trained over a dataset of sentences whose quality is evaluated under such rules. Preliminary results suggest the neural architecture can model such rules to high accuracy.

1 Introduction

People trying to master the art of writing correct sentences in English for the first time often find themselves in the struggle of learning the essential rules to do so. In an effort to assist in this process we studied how to represent a sentence with the purpose of analyzing its quality according to well-known rules. We make use of recent developments in deep learning to achieve a rich representation of single sentences, commonly known as embedding (Palangi et al., 2016). Sentence quality has been traditionally evaluated in the context of machine translation to measure success (Luong et al., 2014; Specia et al., 2009). However, our work measures the quality of a sentence as the end product of the model. Recently, recurrent neural networks (RNNs) with long-short term memories (LSTMs) have been used for word embedding (Ma and Hovy, 2016; Wang et al., 2018), sentence embedding (Vinyals et al., 2015; Chen et al., 2018), and paragraph or document embedding (Li et al., 2015; Zhao et al., 2018) with very good results. These works have demonstrated the utility and robustness of finding discriminative vectorial representations that preserve context and major structures of words and bodies of text. Thus, we decided to study their utility in modeling and assessing the overall quality of individual sentences. A very early version of this work was presented at a workshop in NeurIPS 2018 (Rivas and Pagalla, 2018).

2 Methodology and Results

2.1 Rules

There are well-known rules that make a sentence *a good sentence*. We focused on the following five:

1. **Subjects.** The subject must be the main character, not actions expressed as abstract nouns.
2. **Verbs.** The important actions in the sentence should be verbs, not abstract nouns.
3. **Introductory Phrases.** Introductory phrases in a sentence (if any) should follow rules 1 and 2 and should not be too long; around five words is acceptable.
4. **Nouns.** Strings of consecutive nouns (three or more) should be avoided to preserve sentence clarity.
5. **Conciseness.** Words that mean little or nothing, words that repeat the meaning of other words, words implied by other words, should all be avoided.

2.2 Data

These rules are nearly impossible to define in a precise manner without falling into endless exceptions due to the natural complexities that the English language produces having a context-dependent grammar.

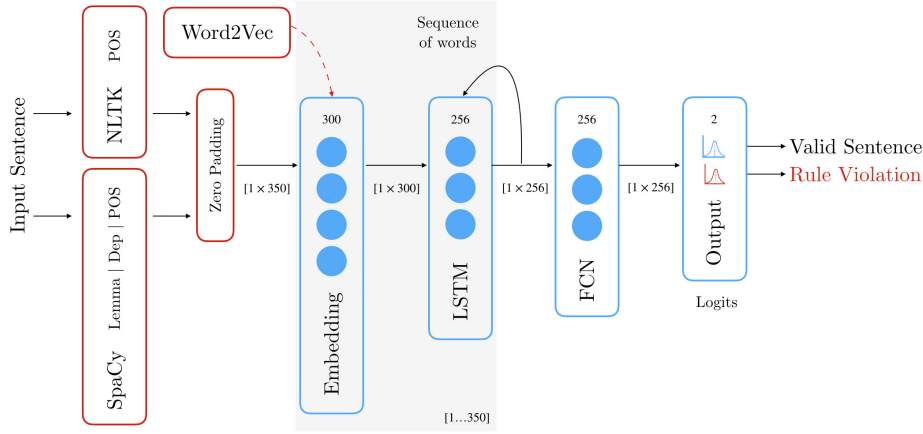


Figure 1: Deep architecture for sentence quality assessment.

Metric	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Avg.
MAE	0.097 ± 0.01	0.109 ± 0.00	0.143 ± 0.01	0.089 ± 0.01	0.229 ± 0.01	0.133
ACC	90.45 ± 0.01	89.23 ± 0.00	85.87 ± 0.01	91.23 ± 0.01	77.43 ± 0.01	86.84

Table 1: Results for each rule. Cross-validated mean absolute error and accuracy.

Thus, we approached the problem by creating a dataset labeled by English experts evaluating a sentence using the five rules above. The data comes from college students writing after the removal of all personal identification data. The average length of the sentences in the dataset is 110 words, and the vocabulary size is 73140.

2.3 Architecture

The input to our model is a sentence of arbitrary length. Following a common practice (Bocklisch et al., 2017), we use the SpaCy NLP library for part-of-speech (POS)-tagging, word dependency analysis, and word lemma identification. Similarly, we use the NLTK library for POS-tagging (Bird and Loper, 2004). One of the reasons to use both libraries is due to their differences in the methodology for POS-tagging.

Once these features are extracted for a sentence, it follows to zero-pad before embedding. The embedding layer has 300 neurons and it is pre-trained using Word2Vec (Mikolov et al., 2013), which in our experiments increased accuracy by 3.2%.

The embedding layer is followed by an LSTM layer that is used to characterize and preserve any potential rule violations. This layer has 256 neural units and has a sequence length of 350, or 350 layers when the recurrence is unfolded. The RNN is followed by a fully connected network (FCN) with 256 neurons and an output layer with 2 neurons using logit activations. These two neurons will output a probability of the violation or fulfillment of each rule. This architecture is depicted in Figure 1. For every single rule, there is a model following the exact same architecture.

2.4 Results

The overall cross-validated average accuracy is 86.84% and the mean absolute error is 0.133 across all five rules. Table 2.4 shows the results for each individual rule in cross validation. The mean absolute error is calculated as $\frac{\sum_{i=1}^N |y_i - \hat{y}_i|}{N}$ where y_i is the i -th target output and \hat{y}_i is the predicted output of the network. The rule most successfully predicted is the rule of nouns with 91.23% while the worst is the rule of concision with 77.43%. Intuitively, the rule about concision is extremely complicated to model as it needs a larger corpus relating meaning of phrases to a summarized expression of the same meaning. Further work will explore this alternative.

A live demo of this project can be found in <https://wa.reev.us> in which we implemented the models to give sentence-level feedback to people learning to write good sentences.

References

- Steven Bird and Edward Loper. 2004. Nltk: the natural language toolkit. In *Proceedings of the ACL 2004 on Interactive poster and demonstration sessions*, page 31. Association for Computational Linguistics.
- Tom Bocklisch, Joey Faulker, Nick Pawlowski, and Alan Nichol. 2017. Rasa: Open source language understanding and dialogue management. *arXiv preprint arXiv:1712.05181*.
- Qian Chen, Zhen-Hua Ling, and Xiaodan Zhu. 2018. Enhancing sentence embedding with generalized pooling. *arXiv preprint arXiv:1806.09828*.
- Jiwei Li, Minh-Thang Luong, and Dan Jurafsky. 2015. A hierarchical neural autoencoder for paragraphs and documents. *arXiv preprint arXiv:1506.01057*.
- Ngoc-Quang Luong, Laurent Besacier, and Benjamin Lecouteux. 2014. Word confidence estimation and its integration in sentence quality estimation for machine translation. In *Knowledge and Systems Engineering*, pages 85–98. Springer.
- Xuezhe Ma and Eduard Hovy. 2016. End-to-end sequence labeling via bi-directional lstm-cnns-crf. *arXiv preprint arXiv:1603.01354*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Hamid Palangi, Li Deng, Yelong Shen, Jianfeng Gao, Xiaodong He, Jianshu Chen, Xinying Song, and Rabab Ward. 2016. Deep sentence embedding using long short-term memory networks: Analysis and application to information retrieval. *IEEE/ACM Transactions on Audio, Speech and Language Processing (TASLP)*, 24(4):694–707.
- Pablo Rivas and Aishwarya Pagalla. 2018. Rule-based sentence quality modeling and assessment using deep lstm features. In *NeurIPS 2018 Workshop on Causal Learning*.
- Lucia Specia, Marco Turchi, Nicola Cancedda, Marc Dymetman, and Nello Cristianini. 2009. Estimating the sentence-level quality of machine translation systems. In *13th Conference of the European Association for Machine Translation*, pages 28–37.
- Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. 2015. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164.
- Jenq-Haur Wang, Ting-Wei Liu, Xiong Luo, and Long Wang. 2018. An lstm approach to short text sentiment classification with word embeddings. In *Proceedings of the 30th Conference on Computational Linguistics and Speech Processing (ROCLING 2018)*, pages 214–223.
- Yue Zhao, Xiaolong Jin, Yuanzhuo Wang, and Xueqi Cheng. 2018. Document embedding enhanced event detection with hierarchical and supervised attention. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, volume 2, pages 414–419.