

Using Integrated Gradients and Constituency Parse Trees to explain Linguistic Acceptability learnt by BERT

Anmol Nayak, Hari Prasad Timmapathini

ARiSE Labs at Bosch

{Anmol.Nayak, HariPrasad.Timmapathini}@in.bosch.com

Abstract

Linguistic Acceptability is the task of determining whether a sentence is grammatical or ungrammatical. It has applications in several use cases like Question-Answering, Natural Language Generation, Neural Machine Translation, where grammatical correctness is crucial. In this paper we aim to understand the decision-making process of BERT (Devlin et al., 2019) in distinguishing between Linguistically Acceptable sentences (LA) and Linguistically Unacceptable sentences (LUA). We leverage Layer Integrated Gradients Attribution Scores (LIG) to explain the Linguistic Acceptability criteria that are learnt by BERT on the Corpus of Linguistic Acceptability (CoLA) (Warstadt et al., 2018) benchmark dataset. Our experiments on 5 categories of sentences lead to the following interesting findings: 1) LIG for LA are significantly smaller in comparison to LUA, 2) There are specific subtrees of the Constituency Parse Tree (CPT) for LA and LUA which contribute larger LIG, 3) Across the different categories of sentences we observed around 88% to 100% of the Correctly classified sentences had positive LIG, indicating a strong positive relationship to the prediction confidence of the model, and 4) Around 43% of the Misclassified sentences had negative LIG, which we believe can become correctly classified sentences if the LIG are parameterized in the loss function of the model.

1 Introduction

Linguistic acceptability is an important criteria in Natural Language Processing and is one of the tasks in the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2018). With the evolution of language encoders like BERT (which leverages the multi-head self-attention mechanism (Vaswani et al., 2017) in its architecture) that have been a breakthrough in language understanding and achieved state-of-the-art

results, the field of probing these architectures for understanding their behaviours has become important.

While there have been several works on interpreting and understanding the different layers of BERT with respect to lexical, syntactic and semantic behaviours (Jawahar et al., 2019; Lin et al., 2019; Clark et al., 2019; Vashishth et al., 2019; Rogers et al., 2020), the focus on explaining the linguistic acceptability (grammaticality) learnt by BERT has been sparse. Some of the recent works have used probing tasks to understand the model’s knowledge on particular grammatical features (Shi et al., 2016; Ettinger et al., 2016; Tenney et al., 2019), relying on language model probabilities to judge grammatical acceptability on sentences that differ minimally (Marvin and Linzen, 2018; Wilcox et al., 2019), or probing the model’s by training with boolean grammaticality judgement objectives (Linzen et al., 2016; Warstadt et al., 2018; Kann et al., 2019; Warstadt et al., 2019). These methods have made significant progress in uncovering that BERT has indeed learnt various aspects of grammatical knowledge, however their focus has not been on explaining the black box details of how BERT arrives at a grammaticality judgement. Our paper attempts to address this by explaining the model’s linguistic acceptability judgement with LIG and CPT (a type of grammar tree) representations.

Attention mechanism based methods (Bahdanau et al., 2014; Vaswani et al., 2017) provide interpretable understanding of the model’s behaviour, however the attention scores cannot be solely relied upon since a feature could influence the output in multiple ways (for e.g. through memory cells, recurrent states etc. in LSTM networks). Feature attribution methods aim to understand the relationship between the model’s output and the input features. They are helpful in interpreting the black-

Sentence	Category
rebecca saw the play .	CIA ^{LA}
the play saw .	CIA ^{LUA}
i surprised myself .	RAA ^{LA}
i surprised himself .	RAA ^{LUA}
the boy is here .	SVA ^{LA}
the boy are here .	SVA ^{LUA}
michael read the book .	SVO ^{LA}
michael the book read .	SVO ^{LUA}
what did rebecca read ?	WHE ^{LA}
what did rebecca read the book ?	WHE ^{LUA}

Table 1: Sample sentences across the 5 categories from the CoLA Targeted Test Sets.

box details of neural networks and provide insights that can be used to improve model performance. While previous feature attribution methods such as DeepLift (Shrikumar et al., 2016, 2017), Layer-wise relevance propagation (Binder et al., 2016) and LIME (Ribeiro et al., 2016) have provided interesting frameworks, they break at least one of the two axioms that are fundamental for attribution methods, namely *Sensitivity* and *Implementation Invariance* (Sundararajan et al., 2017).

In our paper we have chosen the Integrated Gradients (IG) (Sundararajan et al., 2017) attribution method as it satisfies both the aforementioned axioms. IG is a post-hoc interpretability technique which aggregates the gradients of the input by interpolating in small steps along the straight line between a baseline (typically a vector with all zeros) and the input. A large positive or negative IG score indicates that the feature strongly increases or decreases the network output respectively, while a score close to zero indicates that the feature does not influence the network output. This can also be understood as follows: a positive score indicates that the feature tends to agree with the model’s prediction, while a negative score indicates that the feature tends to disagree with the model’s prediction. LIG are computed as the IG between the model output and a particular layer’s input or output. Our work attempts to answer the following Research Questions:

1. Can LIG of a Constituency Parse Tree (CPT) give insights on LA vs LUA?
2. Can LIG be reliably used to explain the Linguistic Acceptability criteria learnt by BERT?
3. Is there a relationship between LIG and the

prediction confidence of the model?

2 Experiment Setup

CoLA dataset sentences have a boolean acceptability judgement, namely LA and LUA. We have used the fine-tuned BERT-Base-Uncased-CoLA model (12 encoder layers with 12 attention heads) provided by TextAttack (Morris et al., 2020), the Captum PyTorch Interpretability library (Kokhlikyan et al., 2020) for computing LIG and the Stanford CoreNLP toolkit (version 4.2.1) (Manning et al., 2014) for constructing the CPT. Integrated Gradients (IG) across the i^{th} dimension of input x and baseline x' are computed as follows:

$$IG = (x_i - x'_i) \times \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha \times (x - x'))}{\partial x_i} d\alpha$$

Captum library approximates the above integral using the Gauss-Legendre approximation algorithm over 50 uniform steps of $\alpha \in [0, 1]$. The baseline was selected as a 768 dimension zero vector. The attribution score for each word is summed across the dimensions (768 in the case of BERT-Base) and normalized using the Euclidean norm of the scores of all the words in the sentence.

We analyzed 5 different categories of sentences within the CoLA Targeted Test Sets: Causative-Inchoative Alternation (CIA), Reflexive-Antecedent Agreement (RAA), Subject-Verb Agreement (SVA), Subject-Verb-Object (SVO) and Wh-Extraction (WHE). A few sample sentences across the categories can be seen in Table 1.

The primary focus of our experiments relied on the LIG computed between the predicted class logit and the token embedding of the words. Further we also computed LIG heatmaps with respect to the Input (Token + Segment + Position) embedding and across the 12 Encoder layer embeddings of BERT to analyze the LIG characteristics. Figure 1 shows the LIG heatmaps of the top 10 CPT patterns for the LA and LUA in the WHE category. The unique CPT patterns were extracted for the correctly classified sentences of each category, corresponding to which the LIG of each subtree were computed. LIG of a subtree is equal to the sum of the LIG of the words appearing as leaf nodes in the subtree. The results in Table 2, Table 3, Figure 2 and Figure 3 represent the LIG computed between the predicted class logit and the token embedding of the words. For Out-Of-Vocabulary words (OOV), the LIG are summed across its tokenized sub-words.

CPT Pattern	Avg. LIG	Category
(S(NP(NN))(VP(VBD)(NP(DT)(NN)))(.))	0.065	CIA ^{LA}
(S(NP(DT)(NN))(VP(VBD)))(.)	1.351	CIA ^{LUA}
(S(NP(PRP))(VP(VBD)(NP(PRP)))(.))	0.061	RAA ^{LA}
(S(NP(PRP))(VP(VBD)(NP(PRP)))(.)	0.926	RAA ^{LUA}
(S(NP(DT)(NN))(VP(VBZ)(ADV(P(RB)))(.))	0.064	SVA ^{LA}
(S(NP(DT)(NN))(VP(VBP)(ADJP(JJ)))(.))	1.067	SVA ^{LUA}
(S(NP(NN))(VP(VBD)(NP(DT)(NN)))(.))	-0.012	SVO ^{LA}
(S(NP(NP(NN))(NP(DT)(NN)))(VP(VBD)))(.)	1.226	SVO ^{LUA}
(SBARQ(WHNP(WP))(SQ(VBD)(NP(NN))(VP(VB)))(.))	0.205	WHE ^{LA}
(SBARQ(WHNP(WP))(SQ(VBD)(NP(NN))(VP(VB)(NP(DT)(NN)))(.))	1.394	WHE ^{LUA}

Table 2: Average normalized LIG of most frequent CPT patterns on 5 categories of Correctly classified CoLA Targeted Test Sets sentences. Subtrees in bold have the largest LIG in the respective categories.

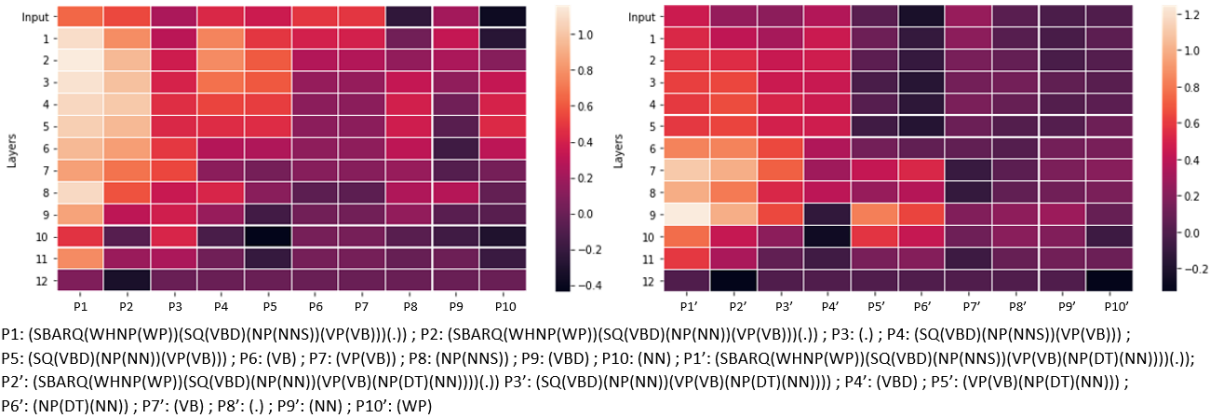


Figure 1: LIG heatmaps of the top 10 scoring CPT patterns ranked in descending order based on averaged LIG across the different BERT layers for LA (Left) and LUA (Right) of the WHE category.

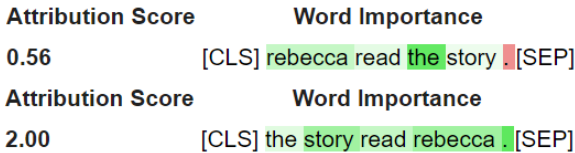


Figure 2: LIG visualization for a LA sentence (Top) and LUA sentence (Bottom) of the SVO category. Green highlighted words contributed strongly towards the model output to be predicted as LA and LUA.

3 LIG for Constituency Parse Tree patterns

CPT is a type of grammar tree which captures the relations between the constituents of a sentence. We believe that analyzing the CPT patterns will give us insights into the grammatical structure of LA and LUA. Computing the LIG for CPT patterns at different subtree levels can give us an indication into the constituents which contribute largely towards making the sentence LA or LUA.

For each of the 5 category of sentences, we extracted all the CPT patterns for correctly classified sentences at every subtree level and picked the most frequent patterns at the root level (Table 2). The subtree patterns in bold are the highest ranking subtrees based on LIG. BERT has been shown to learn surface level features in the early layers, syntactic features in the middle layers and semantic features in the higher layers (Jawahar et al., 2019). Hence, we also wanted to analyse the LIG behaviour across the 12 layers and especially the early to middle layers which are relevant for grammar understanding. Across each of the categories in the LIG heatmaps, it was seen that the top subtree CPT patterns based on token embedding LIG were also dominating across the input and encoder layers of BERT and hence were also found in the top 10 patterns. Further, it can be observed in Figure 1 that there are specific subtrees which dominate more (shades of orange) as compared to others. This characteristic

Category	C	CC	MC	CC ⁺	CC ⁻	MC ⁺	MC ⁻	CC ⁺ %	MC ⁺ %
CIA	182	162	20	144	18	7	13	88.88	35
RAA	144	100	44	100	0	2	42	100	4.54
SVA	676	476	200	441	35	148	52	92.64	74
SVO	500	400	100	362	38	54	46	90.5	54
WHE	520	516	4	465	51	0	4	90.11	0

Table 3: LIG assessment for Correctly classified sentences (CC) and Misclassified (MC) sentences (C: Count, CC⁺: Count of CC having positive LIG, CC⁻: Count of CC having negative LIG, MC⁺: Count of MC having positive LIG, MC⁻: Count of MC having negative LIG, CC⁺%: Percentage of CC⁺ in CC, MC⁺%: Percentage of MC⁺ in MC).




Figure 3: Prediction Probability vs LIG Scatter plots for Correctly classified (Left) and Misclassified (Right) sentences.

is especially useful for debugging LUA as it helps us to understand which phrase contributed largely towards making it unacceptable.

Further, it can be observed in Table 2 that the LIG for LUA are significantly larger than LA. The dominating CPT subtree patterns had a large spike in the LIG for LUA in comparison to LA, indicating that linguistically acceptable patterns were not being adhered. In Figure 2 we can see how the different words in the LA and LUA of the SVO category contributed in varying magnitudes towards the model’s prediction.

4 LIG and Prediction confidence of the model

We investigated to check if there is a relationship between the LIG and the prediction confidence of the model. We found that the range of correctly classified sentences having positive LIG is between 88% to 100% (CC⁺% in Table 3) indicating that whenever the input contributes strongly towards a particular class (whether it is LA or LUA), the model has a higher confidence in making the correct prediction. Around 43% (MC⁻ in Table 3) of the total misclassified sentences had negative LIG which showed that the features disagreed with

the model’s prediction. This behaviour can be observed distinctly in the Figure 3 scatter plots, where we notice that there a large number of points near the top right corner for the correctly classified sentences, and a large number of points near the bottom left corner in the case of misclassified sentences.

We believe that this indication can be used to improve the model’s performance by parameterizing the LIG in the loss function during the later stages of the training process once the model has achieved a reasonable performance (to ensure that the gradients computed are meaningful) and hence serve as a correction mechanism for the model. This aligns with a previous work (Erion et al., 2021) which showed that axiomatic attribution priors improved model performance on many real-world tasks.

5 Conclusion

We have proposed a novel approach for explaining the Linguistic Acceptability criteria learnt by BERT using LIG and CPT patterns. As there is a strong relationship between LIG and the prediction confidence of the model, our future work will focus on parameterizing the LIG in the loss function and observing the model’s performance.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Alexander Binder, Grégoire Montavon, Sebastian Lapuschkin, Klaus-Robert Müller, and Wojciech Samek. 2016. Layer-wise relevance propagation for neural networks with local renormalization layers. In *International Conference on Artificial Neural Networks*, pages 63–71. Springer.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. **What does BERT look at? an analysis of BERT’s attention**. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 276–286, Florence, Italy. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gabriel Erion, Joseph D Janizek, Pascal Sturmfels, Scott M Lundberg, and Su-In Lee. 2021. Improving performance of deep learning models with axiomatic attribution priors and expected gradients. *Nature Machine Intelligence*, pages 1–12.
- Allyson Ettinger, Ahmed Elgohary, and Philip Resnik. 2016. Probing for semantic evidence of composition by means of simple classification tasks. In *Proceedings of the 1st workshop on evaluating vector-space representations for nlp*, pages 134–139.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. What does bert learn about the structure of language? In *ACL 2019-57th Annual Meeting of the Association for Computational Linguistics*.
- Katharina Kann, Alex Warstadt, Adina Williams, and Samuel Bowman. 2019. Verb argument structure alternations in word and sentence embeddings. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 287–297.
- Narine Kokhlikyan, Vivek Miglani, Miguel Martin, Edward Wang, Bilal Alsallakh, Jonathan Reynolds, Alexander Melnikov, Natalia Kliushkina, Carlos Araya, Siqi Yan, et al. 2020. Captum: A unified and generic model interpretability library for pytorch. *arXiv preprint arXiv:2009.07896*.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. **Open sesame: Getting inside BERT’s linguistic knowledge**. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. Assessing the ability of lstms to learn syntax-sensitive dependencies. *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Rose Finkel, Steven Bethard, and David McClosky. 2014. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd annual meeting of the association for computational linguistics: system demonstrations*, pages 55–60.
- Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. *arXiv preprint arXiv:1808.09031*.
- John Morris, Eli Lifland, Jin Yong Yoo, Jake Grigsby, Di Jin, and Yanjun Qi. 2020. Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 119–126.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. ” why should i trust you?” explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. Does string-based neural mt learn source syntax? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1526–1534.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. Learning important features through propagating activation differences. In *International Conference on Machine Learning*, pages 3145–3153. PMLR.
- Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. 2016. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713*.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al.

2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- Shikhar Vashishth, Shyam Upadhyay, Gaurav Singh Tomar, and Manaal Faruqui. 2019. Attention interpretability across nlp tasks. *arXiv preprint arXiv:1909.11218*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*.
- Alex Warstadt, Yu Cao, Ioana Grosu, Wei Peng, Hagen Blix, Yining Nie, Anna Alsop, Shikha Bordia, Haokun Liu, Alicia Parrish, et al. 2019. Investigating bert’s knowledge of language: Five analysis methods with npis. *arXiv preprint arXiv:1909.02597*.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2018. Neural network acceptability judgments. *arXiv preprint arXiv:1805.12471*.
- Ethan Wilcox, Peng Qian, Richard Futrell, Miguel Ballesteros, and Roger Levy. 2019. Structural supervision improves learning of non-local grammatical dependencies. *arXiv preprint arXiv:1903.00943*.