

Inserting Information Bottlenecks for Attribution in Transformers

Zhiying Jiang, Raphael Tang, Ji Xin and Jimmy Lin

David R. Cheriton School of Computer Science
University of Waterloo

{zhiying.jiang, r33tang, ji.xin, jimmylin}@uwaterloo.ca

Abstract

Pretrained transformers achieve the state of the art across tasks in natural language processing, motivating researchers to investigate their inner mechanisms. One common direction is to understand what features are important for prediction. In this paper, we apply information bottlenecks to analyze the attribution of each feature for prediction on a black-box model. We use BERT as the example and evaluate our approach both quantitatively and qualitatively. We show the effectiveness of our method in terms of attribution and the ability to provide insight into how information flows through layers. We demonstrate that our technique outperforms two competitive methods in degradation tests on four datasets. Code is available at <https://github.com/bazingagin/IBA>.

1 Introduction

Increasingly prominent is the urge to interpret deep neural networks, with the success of these black-box models remaining vastly inexplicable both theoretically and empirically. Within natural language processing (NLP), this desire is particularly true for the pretrained transformer, which has witnessed an influx of literature on interpretability analysis. Such papers include visualizing transformer attention mechanisms (Kovaleva et al., 2019), probing the geometry of transformer representations (Hewitt and Manning, 2019), and explaining the span predictions of question answering models (van Aken et al., 2019).

In this paper, we focus on prediction attribution methods. That is, we ask, “Which hidden features contribute the most toward a prediction?” To resolve this question, a number of methods (Selvaraju et al., 2017; Smilkov et al., 2017) generate attribution scores for features, which provide a human-understandable “explanation” of how a particular prediction is made at the instance level. Specif-

ically, given an instance, these methods assign a numerical score for each hidden feature denoting its relevance toward the prediction.

Previous papers have demonstrated that gradient-based methods fail to capture all the information associated with the correct prediction (Li et al., 2016). To address this weakness, Schulz et al. (2020) insert information bottlenecks (Tishby et al., 2000) for attribution, attaining both stronger empirical performance and a theoretical upper bound on the information used. Additionally, mutual information is unconstrained by model and task (Guan et al., 2019). Thus, we adopt information bottlenecks for attribution (IBA) to interpret transformer models at the instance level. We apply IBA to BERT (Devlin et al., 2019) across five datasets in sentiment analysis, textual entailment, and document classification. We show both qualitatively and quantitatively that the method capably captures information in the model’s token-level features, as well as insight into cross-layer behavior.

Our contributions are as follows: First, we are the first to apply information bottlenecks (IB) for attribution to explain transformers. Second, we conduct quantitative analysis to investigate the accuracy of our method compared to other interpretability techniques. Finally, we examine the consistency of our method across layers in a case study. Across four datasets, our technique outperforms integrated gradients (IG) and local interpretable model-agnostic explanations (LIME), two widely adopted prediction attribution approaches.

2 Related Work

In terms of scope, interpretability methods can be categorized as model specific or model agnostic. Model-specific methods interpret only one family of models, whereas model-agnostic techniques aim for wide applicability across many families

of parametric models. We can roughly separate model-agnostic methods into three categories: (1) gradient-based ones (Li et al., 2016; Fong and Vedaldi, 2017; Sundararajan et al., 2017); (2) probing (Ribeiro et al., 2016; Lundberg and Lee, 2017; Tenney et al., 2019; Clark et al., 2019; Liu et al., 2019); (3) information-theoretical methods (Bang et al., 2019; Guan et al., 2019; Schulz et al., 2020; Pimentel et al., 2020).

Gradient-based methods are, however, limited to models with differentiable neural activations. They also fail to capture all the information associated with the correct prediction (Li et al., 2016). Although probing methods provide detailed insight into specific models, they fail to capture inner mechanisms like how information flows through the network (Guan et al., 2019). Information-theoretic methods, in contrast, provide consistent and flexible explanations, as we show in this paper.

Guan et al. (2019) use mutual information to interpret NLP models across different tokens, layers, and neurons, but they lack a quantitative evaluation. Bang et al. (2019) also propose a model-agnostic interpretable model using IB; however, they limit the information through the network by sampling a given number of words at the beginning, which restricts the explanation to neurons only. Our method is inspired by Schulz et al. (2020), who use IBA in image classification.

3 Method

The idea of IBA is to restrict the information flowing through the network for every single instance, such that only the most useful information is kept. Concretely, given an input $\mathbf{X} \in \mathbb{R}^N$ and output $\mathbf{Y} \in \mathbb{R}^M$, an information bottleneck is an intermediate representation \mathbf{T} that maximizes the following function:

$$I(\mathbf{Y}; \mathbf{T}) - \beta \cdot I(\mathbf{X}; \mathbf{T}), \quad (1)$$

where I denotes mutual information and β controls the trade-off between reconstruction $I(\mathbf{Y}; \mathbf{T})$ and information restriction $I(\mathbf{X}; \mathbf{T})$. The larger the β , the narrower the bottleneck, i.e., less information is allowed to flow through the network.

We insert the IB after a given layer l in a pre-trained deep neural network. In this case, $\mathbf{X} = f_l(\mathbf{H})$ represents the chosen layer’s output, where \mathbf{H} is the input of the layer. We restrict information flow by injecting noise into the original input:

$$\mathbf{T} = \boldsymbol{\mu} \odot \mathbf{X} + (\mathbf{1} - \boldsymbol{\mu}) \odot \boldsymbol{\epsilon}, \quad (2)$$

where \odot denotes element-wise multiplication, $\boldsymbol{\epsilon}$ the injected noise, \mathbf{X} the latent representation of the chosen layer, $\mathbf{1}$ the all-one vector, and $\boldsymbol{\mu} \in \mathbb{R}^N$ the weight balancing signal and noise. For every dimension i , $\mu_i \in [0, 1]$, meaning that when $\mu_i = 1$, there is no noise injected into the original representation. To simplify the training process, we set $\mu_i = \sigma(\alpha_i)$, where σ is the sigmoid function and $\boldsymbol{\alpha}$ is a learnable parameter vector. In the extreme case, where all the information in \mathbf{T} is replaced with noise ($\mathbf{T} = \boldsymbol{\epsilon}$), it’s desirable to keep $\boldsymbol{\epsilon}$ the same mean and variance as \mathbf{X} in order to preserve the magnitude of the input to the following layer. Thus, we have $\boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\mu}_{\mathbf{X}}, \sigma_{\mathbf{X}}^2)$.

After obtaining \mathbf{T} , we evaluate how much information \mathbf{T} still contains about \mathbf{X} , which is defined as their mutual information:

$$I(\mathbf{X}; \mathbf{T}) = \mathbb{E}_{\mathbf{X}}[D_{KL}[P(\mathbf{T}|\mathbf{X})||P(\mathbf{T})]], \quad (3)$$

where D_{KL} means Kullback–Leibler (KL) divergence, $P(\mathbf{T}|\mathbf{X})$ and $P(\mathbf{T})$ represent their probability distributions. While $P(\mathbf{T}|\mathbf{X})$ can be sampled empirically, $P(\mathbf{T})$ has no analytical solution since it requires integrating over the feature map $P(\mathbf{T}) = \int P(\mathbf{T}|\mathbf{X})P(\mathbf{X})d\mathbf{X}$. As is standard, we use the variational approximation $Q(\mathbf{T}) = \mathcal{N}(\boldsymbol{\mu}_{\mathbf{X}}, \sigma_{\mathbf{X}}^2)$ to substitute $P(\mathbf{T})$, assuming every dimension of \mathbf{T} is independent and normally distributed. Even though the independence assumption does not hold in general, it only overestimates the mutual information, giving a nice upper bound of mutual information between \mathbf{X} and \mathbf{T} :

$$I(\mathbf{X}; \mathbf{T}) = \mathbb{E}_{\mathbf{X}}[D_{KL}[P(\mathbf{T}|\mathbf{X})||Q(\mathbf{T})]] - D_{KL}[Q(\mathbf{T})||P(\mathbf{T})] \quad (4a)$$

$$I(\mathbf{X}; \mathbf{T}) \leq \mathbb{E}_{\mathbf{X}}[D_{KL}[P(\mathbf{T}|\mathbf{X})||Q(\mathbf{T})]]. \quad (4b)$$

The complete derivation of Equation 4b is in Appendix A. Since we expect $I(\mathbf{X}, \mathbf{T})$ to be small and mutual information to be always nonnegative, the upper bound is a desired property.

Intuitively, the purpose of maximizing $I(\mathbf{Y}; \mathbf{T})$ is to make accurate predictions. Therefore, instead of directly maximizing $I(\mathbf{Y}; \mathbf{T})$, we minimize the loss function for the original task, e.g., the cross entropy \mathcal{L}_{CE} for classification problems after inserting the information bottleneck.

Combining the above two parts, our final loss function \mathcal{L} is

$$\mathcal{L} = D_{KL}[P(\mathbf{T}|\mathbf{X})||Q(\mathbf{T})] + \beta \cdot \mathcal{L}_{CE}. \quad (5)$$

	IMDB	MNLI Matched	MNLI Mismatched	AG News	RTE
Random	0.011	0.106	0.106	0.008	0.012
LIME	0.038	0.244	0.260	0.033	0.014
IG	0.090	0.226	0.233	0.036	0.043
IBA	0.229	0.374	0.367	0.029	0.059

Table 1: Absolute probability drop for the target class after the top 11% most important tokens removed. The larger the score, the more effective the method.

Note that we negate the sign for minimization. The β hyperparameter controls the relative importance between the two loss components. After the optimization process, we obtain for every instance a compressed representation \mathbf{T} .

We then calculate $D_{KL}[P(\mathbf{T}|\mathbf{X})||Q(\mathbf{T})]$, indicating how much information is still kept in \mathbf{T} about \mathbf{X} , which suggests the contribution of each token and feature. To generate the attribution map, we sum over the feature–token axis, obtaining the attribution score of each token.

Overall, we try to learn a compressed hidden representation \mathbf{T} that has just enough information about the input \mathbf{X} to predict the output \mathbf{Y} . This compression is done by adding noise, which removes the least relevant feature-level information, with μ controlling how much to remove.

4 Experiments

Through experimentation, we analyze IBA both quantitatively and qualitatively to understand how it interprets deep neural network across layers.

4.1 Experimental Setting

We compare our method on BERT with two other representative model-agnostic instance-level methods—LIME (Ribeiro et al., 2016), which explores interpretable models for approximation and explanation, and integrated gradients (IG) (Sundararajan et al., 2017), a variation on computing the gradients of the predicted output with respect to input features. For a simple baseline, we also compare with “random,” whose attribution scores are assigned randomly to tokens. On each dataset, we fine-tune BERT and apply these interpretability techniques to the model. We note the test accuracy and generate an attribution score for each token. Details of all parameters are attached in Appendix D.

There is no consensus on how to evaluate interpretability methods quantitatively (Molnar, 2019). LIME’s simulated evaluation leverages the ground

truth of already interpretable models like decision trees, but the ground truth is unavailable for black-box models like neural networks. Therefore, we follow Ancona et al. (2018) and Hooker et al. (2018) and carry out a *degradation test* on IMDB (Maas et al., 2011), AG News (Gulli, 2004), MNLI (Williams et al., 2018), and RTE (Wang et al., 2018), covering sentiment analysis, natural language inference, and text classification.

The degradation test has the following steps:

1. Generate attribution scores s for each interpretability method f : $s = f(\mathcal{M}, x, y)$, where x is the test instance, y is the target label, and \mathcal{M} is the model.
2. Sort tokens by their attribution score in descending order.
3. Remove top k tokens to obtain x' , the degraded instance; k can be preset.
4. Test the target class probability $p(y|x')$ with the original model on the degraded instance.
5. Repeat steps 3 and 4 until all tokens removed.

For the final visualization, we average all test instances at each degradation step to compute $\bar{p}(y|x')$. Then, we normalize the degradation test result $\bar{p}(y|x')$ to $[0, 1]$ using the normalized probability drop $\bar{d} = \frac{\bar{p}(y|x') - m}{o - m}$, where o means the original probability on the nondegraded instance, and m means the minimum of the *fully* degraded instance’s probability across all interpretability models. In this way, the normalized probability drop \bar{d} will be independent of the original model quality and easily comparable across models. Note that, for IBA, we perform the degradation test on the original model, not the one with the inserted bottleneck. Thus, a large β does not directly cause the probability to drop. An effective attribution map can find the most important tokens, which means $\bar{p}(y|x')$ after the degradation step will drop substantially.

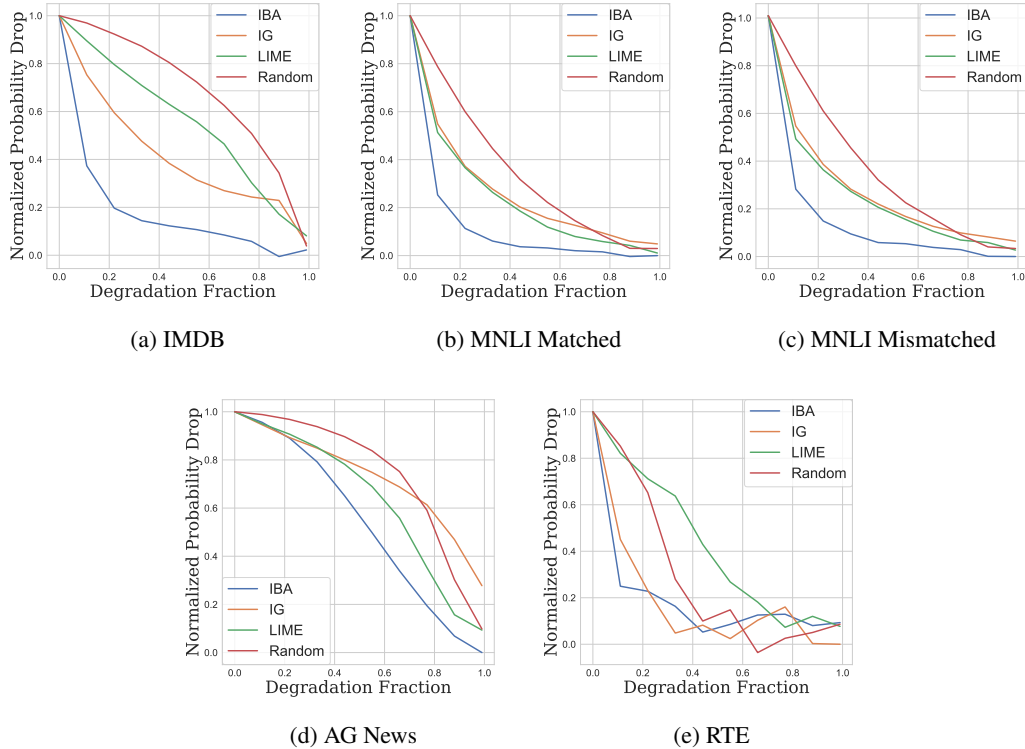


Figure 1: Degradation test results comparing IBA, IG, LIME, and random.

4.2 Results and Analysis

Overall, the results show that our method better identifies the most important tokens compared to other model-agnostic interpretability methods.

Quantitative Analysis. Table 1 shows the absolute probability drop $\|\bar{p}(y|x) - o\|$ with the first 11% of the important tokens removed. We further plot the normalized probability drop after each percentage of the important tokens is removed, as shown in Figure 1, indicating how much important information is lost for prediction: the steeper the slope, the better the ability to capture important tokens. For this experiment, we insert the information bottleneck after layer 9, and we see that removing important tokens that are identified by our method deteriorates the probability the most on IMDB and MNLI Matched/Mismatched.

Of course, choosing the right layer to insert the information bottleneck is crucial to the result. It also indicates which layer encodes the most meaningful information for prediction. To investigate differences in inserting information bottlenecks after different layers, we carry the degradation test on 1000 random test samples across layers on IMDB, as shown in Figure 2a—see Appendix B for all 12 layers. Insertion after layers 1, 8, and 9 generates

more meaningful attribution scores. At layer 1, the tokens remain distinct (i.e., representations have not been aggregated), and it is likely that the latent representation \mathbf{T} is essentially capturing per-token sentiment values. The big drop of \bar{d} after layers 8 and 9, on the other hand, is interesting. Recently, [Xin et al. \(2020\)](#) examined early exit mechanisms in BERT and found that halting inference at layers 8 or 9 produces results not much worse than full inference, which suggests that an abundance of information is encoded in those layers.

Another important parameter is β , which controls the trade-off between restricting the information flow and achieving greater accuracy. A smaller β allows more information through, and an extremely small β has the same effect of using \mathbf{X} as the attribution map. As Figure 2b shows, when $\beta \leq 1e - 6$, the degradation curve is similar to the one using \mathbf{X} only. Appendix C shows the effects of different β on a specific example.

Qualitative Analysis. The first plot in Figure 3 shows the before and after comparison of IB insertion, with positive tokens highlighted. The second and third plots visualize attribution maps for instances across layers. Consistent with our quantitative analysis in Figure 2a, these plots demonstrate




Figure 2: Analysis of different layers and different β .




Figure 3: Illustrations from left to right are as follows: The before and after comparison of inserting an information bottleneck after layer 6; attribution for an IMDB example with the positive label; attribution for an MNLI example with the contradiction label.

that, for a fully fine-tuned BERT, layers 8 and 9 seem to encode the most important information for the prediction. For example, in the IMDB instance, *liked* and *intrigued* have the highest attribution scores for the prediction of positive sentiment across most layers—see layer 9 in particular. In the MNLI example, *never* is mostly highlighted starting from layer 7 to predict “contradiction.”

5 Conclusion

In this paper, we adopt an information-bottleneck-based approach to analyze attribution for transformers. Our method outperforms two widely used attribution methods across four datasets in sentiment analysis, document classification, and textual entailment. We also analyze the information across layers both quantitatively and qualitatively.

Acknowledgments

This research was supported in part by the Canada First Research Excellence Fund and the Natural Sciences and Engineering Research Council (NSERC) of Canada.

References

- Betty van Aken, Benjamin Winter, Alexander Löser, and Felix A. Gers. 2019. How does BERT answer questions? A layer-wise analysis of transformer representations. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*.
- Marco Ancona, Enea Ceolini, Cengiz Öztireli, and Markus Gross. 2018. Towards better understanding of gradient-based attribution methods for deep neural networks. In *International Conference on Learning Representations*.
- Seojin Bang, Pengtao Xie, Heewook Lee, Wei Wu, and Eric Xing. 2019. Explaining a black-box using deep variational information bottleneck approach. *arXiv:1902.06918*.
- 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*.
- 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*.
- Ruth C. Fong and Andrea Vedaldi. 2017. Interpretable explanations of black boxes by meaningful perturbation. In *Proceedings of the IEEE International Conference on Computer Vision*.
- Chaoyu Guan, Xiting Wang, Quanshi Zhang, Runjin Chen, Di He, and Xing Xie. 2019. Towards a deep and unified understanding of deep neural models in NLP. In *International Conference on Machine Learning*.
- Antonio Gulli. 2004. AGNews. http://groups.di.unipi.it/~gulli/AG_corpus_of_news_articles.html.
- John Hewitt and Christopher D. Manning. 2019. A structural probe for finding syntax in word representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. 2018. Evaluating feature importance estimates. *arXiv:1806.10758*.
- Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. 2019. Revealing the dark secrets of BERT. 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)*.
- 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*.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E Peters, and Noah A. Smith. 2019. Linguistic knowledge and transferability of contextual representations. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*.
- Christoph Molnar. 2019. *Interpretable Machine Learning*. <https://christophm.github.io/interpretable-ml-book/>.
- Tiago Pimentel, Josef Valvoda, Rowan Hall Maudslay, Ran Zmigrod, Adina Williams, and Ryan Cotterell. 2020. Information-theoretic probing for linguistic structure. *arXiv:2004.03061*.
- 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*.
- Karl Schulz, Leon Sixt, Federico Tombari, and Tim Landgraf. 2020. Restricting the flow: Information bottlenecks for attribution. In *International Conference on Learning Representations*.
- Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2017. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*.
- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Viégas, and Martin Wattenberg. 2017. SmoothGrad: removing noise by adding noise. *arXiv:1706.03825*.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R. Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel Bowman, Dipanjan Das, and Ellie Pavlick. 2019. What do you learn from context? Probing for sentence structure in contextualized word representations. In *International Conference on Learning Representations*.
- Naftali Tishby, Fernando C. Pereira, and William Bialek. 2000. The information bottleneck method. *arXiv:physics/0004057*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*.
- Ji Xin, Raphael Tang, Jaejun Lee, Yaoliang Yu, and Jimmy Lin. 2020. DeeBERT: Dynamic early exiting for accelerating BERT inference. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.

A Proof of Variational Upper Bound

$$\begin{aligned}
I(\mathbf{X}; \mathbf{T}) &= \mathbb{E}_{\mathbf{X}}[D_{KL}[P(\mathbf{T}|\mathbf{X})||P(\mathbf{T})]] \\
&= \int_{\mathbf{X}} p(x) \left(\int_{\mathbf{T}} p(t|x) \log \frac{p(t|x)}{p(t)} dt \right) dx \\
&= \int_{\mathbf{X}} \int_{\mathbf{T}} p(x, t) \log \frac{p(t|x) q(t)}{p(t) q(t)} dt dx \\
&= \int_{\mathbf{X}} \int_{\mathbf{T}} p(x, t) \log \frac{p(t|x)}{q(t)} dt dx \\
&+ \int_{\mathbf{X}} \int_{\mathbf{T}} p(x, t) \log \frac{q(t)}{p(t)} dt dx \\
&= \int_{\mathbf{X}} \int_{\mathbf{T}} p(x, t) \log \frac{p(t|x)}{q(t)} dt dx \\
&+ \int_{\mathbf{T}} p(t) \left(\int_{\mathbf{X}} p(x|t) dx \right) \log \frac{q(t)}{p(t)} dt \\
&= \mathbb{E}_{\mathbf{X}}[D_{KL}[P(\mathbf{T}|\mathbf{X})||Q(\mathbf{T})]] \\
&- D_{KL}[Q(\mathbf{T})||P(\mathbf{T})] \\
&\leq \mathbb{E}_{\mathbf{X}}[D_{KL}[P(\mathbf{T}|\mathbf{X})||Q(\mathbf{T})]]
\end{aligned}$$

B Degradation Test across 12 Layers

Figure 4 shows the complete version of the degradation test across all 12 layers. In general, the earlier we insert the bottleneck, the larger the probability drop is, except for layers 8 and 9, which are the only two layers with steeper slopes than layer 1.




Figure 4: Degradation test results across all layers.

C Visualization of the Effects of β

Figure 5 shows the effects of different β on a specific example. As we can see, when β is as small as 10^{-7} , most information is allowed to flow through the network and thus most parts are highlighted.

In contrast, when β is larger, the representation is more restricted.

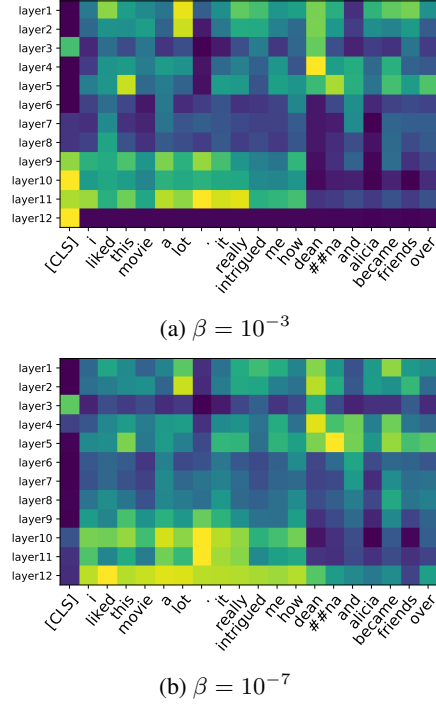


Figure 5: Comparison of BERT attribution maps with different values of β .

D Detailed Parameters and Dataset Information

To keep as much information as possible at the beginning, μ_i should be set close to $1, \forall i$, in which case $\mathbf{T} \approx \mathbf{X}$. So we initialize with $\alpha_i = 5, \forall i$ and therefore $\mu_i \approx 0.993$. In order to stabilize the result, the input of the bottleneck (\mathbf{X}) is duplicated 10 times with different noise added. We set the learning rate to 1 and the number of training steps to 10. We use empirical estimation for $\beta \approx 10 \times \frac{\mathcal{L}_{CE}}{\mathcal{L}_{IB}}$. For IMDB, MNLI Matched/Mismatched, and AGNews, we insert the IB after layer 9 and β is set to 10^{-5} . For RTE, we insert the IB after layer 10 and β is set to 10^{-4} .

We carry out experiments on NVIDIA RTX 2080 Ti GPUs with 11GB VRAM running PyTorch 1.4.0 and CUDA 10.0. A full technical description of our computing environment is released alongside our codebase. For LIME, we set N , the number of permuted samples drawn from the original dataset, to 100 as this reaches the limitation of GPU memory. Similarly, the number of steps of integrated gradients is set to 10 because it is more memory intensive. The average time of running 25000 in-

stances on the described GPU is about 10 hours for IBA, 13 hours for LIME, and 2 hours for IG.

Dataset	Number of Dev/Test
IMDB	25000
MNLI Matched	9815
MNLI Mismatched	9832
AG News	7600
RTE	277

Table 2: Dataset Details.

We use the test sets when the label is provided and use the dev sets otherwise. See Table 2 for details. Note that “IMDB” refers to the sentiment analysis dataset provided by [Maas et al. \(2011\)](#). “MNLI Matched” means that the training set and the test set have the same set of genres while “MNLI Mismatched” means that genres that appear in the test set don’t appear in the training set. Detailed information of the MNLI dataset can be found in [Williams et al. \(2018\)](#).