

Many Faces of Feature Importance: Comparing Built-in and Post-hoc Feature Importance in Text Classification

Vivian Lai and Zheng Cai and Chenhao Tan

Department of Computer Science
University of Colorado Boulder
Boulder, CO

vivian.lai, jon.z.cai, chenhao.tan@colorado.edu

Abstract

Feature importance is commonly used to explain machine predictions. While feature importance can be derived from a machine learning model with a variety of methods, the consistency of feature importance via different methods remains understudied. In this work, we systematically compare feature importance from built-in mechanisms in a model such as attention values and post-hoc methods that approximate model behavior such as LIME. Using text classification as a testbed, we find that 1) no matter which method we use, important features from traditional models such as SVM and XGBoost are more similar with each other, than with deep learning models; 2) post-hoc methods tend to generate more similar important features for two models than built-in methods. We further demonstrate how such similarity varies across instances. Notably, important features do *not* always resemble each other better when two models agree on the predicted label than when they disagree.

1 Introduction

As machine learning models are adopted in societally important tasks such as recidivism prediction and loan approval, explaining machine predictions has become increasingly important (Doshi-Velez and Kim, 2017; Lipton, 2016). Explanations can potentially improve the trustworthiness of algorithmic decisions for decision makers, facilitate model developers in debugging, and even allow regulators to identify biased algorithms.

A popular approach to explaining machine predictions is to identify important features for a particular prediction (Luong et al., 2015; Ribeiro et al., 2016; Lundberg and Lee, 2017). Typically, these explanations assign a value to each feature (usually a word in NLP), and thus enable visualizations such as highlighting top k features.

In general, there are two classes of methods: 1) built-in feature importance that is embedded in the machine learning model such as coefficients in linear models and attention values in attention mechanisms; 2) post-hoc feature importance through credit assignment based on the model such as LIME. It is well recognized that robust evaluation of feature importance is challenging (Jain and Wallace, 2019; Nguyen, 2018, *inter alia*), which is further complicated by different use cases of explanations (e.g., for decision makers vs. for developers). Throughout this work, we refer to machine learning models that learn from data as *models* and methods to obtain local explanations (i.e., feature importance in this work) for a prediction by a model as *methods*.

While prior research tends to focus on the internals of models in designing and evaluating methods of explanations, e.g., how well explanations reflect the original model (Ribeiro et al., 2016), we view feature importance itself as a subject of study, and aim to provide a systematic characterization of important features obtained via different methods for different models. This view is particularly important when explanations are used to support decision making because they are the only exposure to the model for decision makers. It would be desirable that explanations are consistent across different instances. In comparison, debugging represents a distinct use case where developers often know the mechanism of the model beyond explanations. Our view also connects to studying explanation as a *product* in cognitive studies of explanations (Lombrozo, 2012), and is complementary to the model-centric perspective.

Given a wide variety of models and methods to generate feature importance, there are basic open questions such as how similar important features are between models and methods, how important features distribute across instances, and what lin-

One of favorite places to eat on the King W side, simple and relatively quick. I typically always get the chicken burrito and the small is enough for me for dinner. Ingredients are always fresh and watch out for the hot sauce cause it's skull scratching hot. Seating is limited so be prepared to take your burrito outside or you can even eat at Metro Hall Park.

models methods \	SVM (ℓ_2)	XGBoost	LSTM with attention	BERT
built-in	sauce, seating, park, prepared, even, always, can, fresh, quick, favorite	is, can, quick, fresh, at, to, always, even, favorite, and	me, be, relatively, enough, always, fresh, ingredients, prepared, quick, favorite	.., ingredients, relatively, quick, places, enough, dinner, typically, me, i
LIME	the, dinner, be, quick, and, even, you, always, fresh, favorite	you, to, fresh, quick, at, can, even, always, and, favorite	dinner, ingredients, typically, fresh, places, cause, quick, and, favorite, always	one, watch, to, enough, limited, cause, and, fresh, hot, favorite

Table 1: 10 most important features (separated by comma) identified by different methods for different models for the given review. In the interest of space, we only show built-in and LIME here.

guistic properties important features tend to have. We use text classification as a testbed to answer these questions. We consider built-in importance from both traditional models such as linear SVM and neural models with attention mechanisms, as well as post-hoc importance based on LIME and SHAP. Table 1 shows important features for a Yelp review in sentiment classification. Although most approaches consider “fresh” and “favorite” important, there exists significant variation.

We use three text classification tasks to characterize the overall similarity between important features. Our analysis reveals the following insights:

- (Comparison between approaches) Deep learning models generate more different important features from traditional models such as SVM and XGBoost. Post-hoc methods tend to reduce the dissimilarity between models by making important features more similar than the built-in method. Finally, different approaches do not generate more similar important features even if we focus on the most important features (e.g., top one feature).
- (Heterogeneity between instances) Similarity between important features is not always greater when two models agree on the predicted label, and longer instances are less likely to share important features.
- (Distributional properties) Deep models generate more diverse important features with higher entropy, which indicates lower consistency across instances. Post-hoc methods bring the POS distribution closer to background distributions.

In summary, our work systematically compares important features from different methods for different models, and sheds light on how different models/methods induce important features. Our

work takes the first step to understand important features as a product and helps inform the adoption of feature importance for different purposes. Our code is available at <https://github.com/BoulderDS/feature-importance>.

2 Related Work

To provide further background for our work, we summarize current popular approaches to generating and evaluating explanations of machine predictions, with an emphasis on feature importance.

Approaches to generating explanations. A battery of approaches have been recently proposed to explain machine predictions (see Guidotti et al. (2019) for an overview), including example-based approaches that identifies “informative” examples in the training data (e.g., Kim et al., 2016) and rule-based approaches that reduce complex models to simple rules (e.g., Malioutov et al., 2017). Our work focuses on characterizing properties of feature-based approaches. Feature-based approaches tend to identify important features in an instance and enable visualizations with important features highlighted. We discuss several directly related post-hoc methods here and introduce the built-in methods in §3. A popular approach, LIME, fits a sparse linear model to approximate model predictions locally (Ribeiro et al., 2016); Lundberg and Lee (2017) present a unified framework based on Shapley values, which can be computed with different approximation methods for different models. Gradients are popular for identifying important features in deep learning models since these models are usually differentiable (Shrikumar et al., 2017), for instance, Li et al. (2016) uses gradient-based saliency to compare LSTMs with simple recurrent networks.

Definition and evaluation of explanations. De-

spite a myriad of studies on approaches to explaining machine predictions, explanation is a rather overloaded term and evaluating explanations is challenging. [Doshi-Velez and Kim \(2017\)](#) lays out three levels of evaluations: functionally-grounded evaluations based on proxy automatic tasks, human-grounded evaluations with laypersons on proxy tasks, and application-grounded based on expert performance in the end task. In text classification, [Nguyen \(2018\)](#) shows that automatic evaluation based on word deletion moderately correlate with human-grounded evaluations that ask crowdworkers to infer machine predictions based on explanations. However, explanations that help humans infer machine predictions may not actually help humans make better decisions/predictions. In fact, recent studies find that feature-based explanations alone have limited improvement on human performance in detecting deceptive reviews and media biases ([Lai and Tan, 2019](#); [Horne et al., 2019](#)).

In another recent debate, [Jain and Wallace \(2019\)](#) examine attention as an explanation mechanism based on how well attention values correlate with gradient-based feature importance and whether they exclusively lead to the predicted label, and conclude that attention is not explanation. Similarly, [Serrano and Smith \(2019\)](#) show that attention is not a fail-safe indicator for explaining machine predictions based on intermediate representation erasure. However, [Wiegreffe and Pinter \(2019\)](#) argue that attention can be explanation depending on the definition of explanations (e.g., plausibility and faithfulness).

In comparison, we treat feature importance itself as a subject of study and compare different approaches to obtaining feature importance from a model. Instead of providing a normative judgment with respect to what makes good explanations, our goal is to allow decision makers or model developers to make informed decisions based on properties of important features using different models and methods.

3 Approach

In this section, we first formalize the problem of obtaining feature importance and then introduce the models and methods that we consider in this work. Our main contribution is to compare important features identified for a particular instance through different methods for different models.

Feature importance. For any instance t and a machine learning model $m : t \rightarrow y \in \{0, 1\}$, we use method h to obtain feature importance on an interpretable representation of t , $\mathcal{I}_t^{m,h} \in \mathbb{R}^d$, where d is the dimension of the interpretable representation. In the context of text classification, we use unigrams as the interpretable representation. Note that the machine learning model does not necessarily use the interpretable representation. Next, we introduce the models and methods in this work.

Models (m). We include both recent deep learning models for NLP and popular machine learning models that are not based on neural networks. In addition, we make sure that the chosen models have some built-in mechanism for inducing feature importance and describe the built-in feature importance as we introduce the model.¹

- Linear SVM with ℓ_2 (or ℓ_1) regularization. Linear SVM has shown strong performance in text categorization ([Joachims, 1998](#)). The absolute value of coefficients in these models is typically considered a measure of feature importance (e.g., [Ott et al., 2011](#)). We also consider ℓ_1 regularization because ℓ_1 regularization is often used to induce sparsity in the model.
- Gradient boosting tree (XGBoost). XGBoost represents an ensembled tree algorithm that shows strong performance in competitions ([Chen and Guestrin, 2016](#)). We use the default option in XGBoost to measure feature importance with the average training loss gained when using a feature for splitting.
- LSTM with attention (often shortened as LSTM in this work). Attention is a commonly used technique in deep learning models for NLP ([Bahdanau et al., 2015](#)). The intuition is to assign a weight to each token before aggregating into the final prediction (or decoding in machine translation). We use the dot product formulation in [Luong et al. \(2015\)](#). The weight on each token has been commonly used to visualize the importance of each token. To compare with the previous bag-of-words models, we use the average weight of each type (unique token) in this work to measure feature importance.
- BERT. BERT represents an example architecture based on Transformers, which could show different behavior from LSTM-style recurrent

¹For instance, we do not consider LSTM as a model here due to the lack of commonly-accepted built-in mechanisms.

networks (Devlin et al., 2019; Vaswani et al., 2017; Wolf, 2019). It also achieves state-of-the-art performance in many NLP tasks. Similar to LSTM with attention, we use the average attention values of 12 heads used by the first token at the final layer (the representation passed to fully connected layers) to measure feature importance for BERT.² Since BERT uses a subword tokenizer, for each word, we aggregate the attention on related subparts. BERT also requires special processing due to the length constraint; please refer to the supplementary material for details. As a result, we focus on presenting LSTM with attention in the main paper for ease of understanding.

Methods (h). For each model, in addition to the built-in feature importance that we described above, we consider the following two popular methods for extracting post-hoc feature importance (see the supplementary material for details of using the post-hoc methods).

- **LIME** (Ribeiro et al., 2016). LIME generates post-hoc explanations by fitting a local sparse linear model to approximate model predictions. As a result, the explanations are sparse.
- **SHAP** (Lundberg and Lee, 2017). SHAP unifies several interpretations of feature importance through Shapley values. The main intuition is to account the importance of a feature by examining the change in prediction outcomes for all the combinations of other features. Lundberg and Lee (2017) propose various approaches to approximate the computation for different classes of models (including gradient-based methods for deep models).

Note that feature importances obtained via all approaches are all local, because the top features are conditioned on an instance (i.e., words present in an instance) even for the built-in method for SVM and XGBoost.

Comparing feature importance. Given $\mathcal{I}_t^{m,h}$ and $\mathcal{I}_t^{m',h'}$, we use Jaccard similarity based on the top k features with the greatest absolute feature importance, $\frac{|\text{TopK}(\mathcal{I}_t^{m,h}) \cap \text{TopK}(\mathcal{I}_t^{m',h'})|}{|\text{TopK}(\mathcal{I}_t^{m,h}) \cup \text{TopK}(\mathcal{I}_t^{m',h'})|}$, as our main

²We also tried to use the max of 12 heads and previous layers, and the average of the final layer is more similar to SVM (ℓ_2) than the average of first layer. Results are in the supplementary material. Vig (2019) show that attention in BERT tends to be on first words, neighboring words, and even separators. The complex choices for BERT further motivate our work to view feature importance as a subject of study.

similarity metric for two reasons. First, the most typical way to use feature importance for interpretation purposes is to show the most important features (Lai and Tan, 2019; Ribeiro et al., 2016; Horne et al., 2019). Second, some models and methods inherently generate sparse feature importance, so most feature importance values are 0.

It is useful to discuss the implication of similarity before we proceed. On the one hand, it is possible that different models/methods identify the same set of important features (high similarity) and the performance difference in prediction is due to how different models weigh these important features. If this were true, the choice of model/method would have mattered little for visualizing important features. On the other hand, a low similarity poses challenges for choosing which model/method to use for displaying important features. In that case, this work aims to develop an understanding of how the similarity varies depending on models and methods, instances, and features. We leave it to future work for examining the impact on human interaction with feature importance. Low similarity may enable model developers to understand the differences between models, but may lead to challenges for decision makers to get a consistent picture of what the model relies on.

4 Experimental Setup and Hypotheses

Our goal is to characterize the similarities and differences between feature importances obtained with different methods and different models. In this section, we first present our experimental setup and then formulate our hypotheses.

Experimental setup. We consider the following three text classification tasks in this work. We choose to focus on classification because classification is the most common scenario used for examining feature importance and the associated human interpretation (e.g., Jain and Wallace, 2019).

- **Yelp** (Yelp, 2019). We set up a binary classification task to predict whether a review is positive (rating ≥ 4) or negative (rating ≤ 2). As the original dataset is huge, we subsample 12,000 reviews for this work.
- **SST** (Socher et al., 2013). It is a sentence-level sentiment classification task and represents a common benchmark. We only consider the binary setup here.
- **Deception detection** (Ott et al., 2013, 2011).

Model	Yelp	SST	Deception
SVM (ℓ_2)	92.3	80.8	86.3
SVM (ℓ_1)	91.5	79.2	84.4
XGBoost	88.8	75.9	83.4
LSTM w/ attention	93.9	82.6	88.4
BERT	95.5	92.2	90.9

Table 2: Accuracy on the test set.

This dataset was created by extracting genuine reviews from TripAdvisor and collecting deceptive reviews using Turkers. It is relatively small with 1,200 reviews and represents a distinct task from sentiment classification.

For all the tasks, we use 20% of the dataset as the test set. For SVM and XGBoost, we use cross validation on the other 80% to tune hyperparameters. For LSTM with attention and BERT, we use 10% of the dataset as a validation set, and choose the best hyperparameters based on the validation performance. We use spaCy to tokenize and obtain part-of-speech tags for all the datasets (Honnibal and Montani, 2017). Table 2 shows the accuracy on the test set and our results are comparable to prior work. Not surprisingly, BERT achieves the best performance in all three tasks. For important features, we use $k \leq 10$ for Yelp and deception detection, and $k \leq 5$ for SST as it is a sentence-level task. See supplementary materials for details of preprocessing, learning, and dataset statistics.

Hypotheses. We aim to examine the following three research questions in this work: 1) How similar are important features between models and methods? 2) What factors relate to the heterogeneity across instances? 3) What words tend to be chosen as important features?

Overall similarity. Here we focus on discussing comparative hypotheses, but we would like to note that it is important to understand to what extent important features are similar across models (i.e., the value of similarity score). First, as deep learning models and XGBoost are nonlinear, we hypothesize that built-in feature importance is more similar between SVM (ℓ_1) and SVM (ℓ_2) than other model pairs (**H1a**). Second, LIME generates more similar important features to SHAP than to built-in feature importance because both LIME and SHAP make additive assumptions, while built-in feature importance is based on drastically different models (**H1b**). It also follows that post-hoc explanations of different models show higher similarity than built-in ex-




Figure 1: Jaccard similarity between the top 10 features of different models based on built-in feature importance on Yelp. The similarity is the greatest between SVM (ℓ_2) and SVM (ℓ_1), while LSTM with attention and BERT pay attention to quite different features from other models.

planations across models. Third, the similarity with small k is higher (**H1c**) because hopefully, all models and methods agree what the most important features are.

Heterogeneity between instances. Given a pair of (model, method) combinations, our second question is concerned with how instance-level properties affect the similarity in important features between different combinations. We hypothesize that 1) when two models agree on the predicted label, the similarity between important features is greater (**H2a**); 2) longer instances are less likely to share similar important features (**H2b**). 3) instances with higher type-token ratio,³ which might be more complex, are less likely to share similar important features (**H2c**).

Distribution of important features. Finally, we examine what words tend to be chosen as important features. This question certainly depends on the nature of the task, but we would like to understand how consistent different models and methods are. We hypothesize that 1) deep learning models generate more diverse important features (**H3a**); 2) adjectives are more important in sentiment classification, while pronouns are more important in deception detection as shown in prior work (**H3b**).

5 Similarity between Instance-level Feature Importance

We start by examining the overall similarity between different models using different methods. In a nutshell, we compute the average Jaccard similarity of top k features for each pair of (m, h) and

³Type-token ratio is defined as the number of unique tokens divided by the number of tokens.




Figure 2: Similarity comparison between models with the same method. x -axis represents the number of important features that we consider, while y -axis shows the Jaccard similarity. *Error bars represent standard error throughout the paper.* The top row compares three pairs of models using the built-in method, while the second row compares three methods on SVM and LSTM with attention (LSTM in figure legends always refers to LSTM with attention in this work). The random line is derived using the average similarity between two random samples of k features from 100 draws.

(m', h') . To facilitate effective comparisons, we first fix the method and compare the similarity of different models, and then fix the model and compare the similarity of different methods. Figure 1 shows the similarity between different models using the built-in feature importance for the top 10 features in Yelp ($k = 10$). Consistent with **H1a**, SVM (ℓ_2) and SVM (ℓ_1) are very similar to each other, and LSTM with attention and BERT clearly lead to quite different top 10 features from the other models. As the number of important features (k) can be useful for evaluating the overall trend, we thus focus on line plots as in Figure 2 in the rest of the paper. This heatmap visualization represents a snapshot for $k = 10$ using the built-in method. Also, we only include SVM (ℓ_2) in the main paper for ease of visualization and sometimes refer to it in the rest of the paper as SVM.

No matter which method we use, important features from SVM and XGBoost are more similar with each other, than with deep learning models (Figure 2). First, we compare the similarity of feature importance between different models us-

ing the same method. Using the built-in method (first row in Figure 2), the solid line (SVM x XG-Boost) is always above the other lines, usually by a significant margin, suggesting that deep learning models such as LSTM with attention are less similar to traditional models. In fact, the similarity between XGBoost and LSTM with attention is lower than random samples for $k = 1, 2$ in SST. Similar results also hold for BERT (see supplementary materials). Another interesting observation is that post-hoc methods tend to generate greater similarity than built-in methods, especially for LIME (the dashed line (LIME) is always above the solid line (built-in) in the second row of Figure 2). This is likely because LIME only depends on the model behavior (i.e., what the model predicts) and does not account for how the model works.

The similarity between important features from different methods tends to be lower for LSTM with attention (Figure 3). Second, we compare the similarity of feature importance derived from the same model with different methods. For deep learning models such as LSTM with at-




Figure 3: Similarity comparison between methods using the same model. The similarity between different methods based on LSTM with attention is generally lower than other methods. Similar results hold for BERT (see the supplementary material).




Figure 4: Similarity between SVM (ℓ_2) and LSTM with attention with different methods grouped by whether these two models agree on the predicted label. The similarity is not always greater when they agree on the predicted labels than when they disagree.




Figure 5: In most cases, the similarity between feature importance is negatively correlated with length. Here we only show the comparison between different methods based on the same model. Similar results hold for comparison between different models using the same method. For ease of comparison, the gray line marks the value 0. Generally as k grows, relationship becomes even more negatively correlated.

tention, the similarity between feature importance generated by different methods is the lowest, especially comparing LIME with SHAP. Notably, the results are much more cluttered in deception detection. Contrary to **H1b**, we do not observe that LIME is more similar to SHAP than built-in. The order seems to depend on both the task and the model: even within SST, the similarity between built-in and LIME can rank as third, second, or

first. In other words, post-hoc methods generate more similar important features when we compare different models, but that is not the case when we fix the model. It is reassuring that that similarity between any pairs is above random, with a sizable margin in most cases (BERT on SHAP is an exception; see supplementary materials).

Relation with k . As the relative order between different approaches can change with k , we have




Figure 6: The entropy of important features. LSTM with attention generates more diverse important features than SVM and XGBoost.




Figure 7: Part-of-speech tag distribution with the built-in method. “Background” shows the distribution of all words in the test set. LSTM with attention puts a strong emphasis on nouns in deception detection, but is not necessarily more different from the background than other models.

so far only focused on relatively consistent patterns over k and classification tasks. Contrary to **H1c**, the similarity between most approaches is not drastically greater for small k , which suggests that different approaches may not even agree on the most important features. In fact, there is no consistent trend as k grows: similarity mostly *increases* in SST (while our hypothesis is that it decreases), increases or stays level in Yelp, and shows varying trends in deception detection.

6 Heterogeneity between Instances

Given the overall low similarity between different methods/models, we next investigate how the similarity may vary across instances.

The similarity between models is not always greater when two models agree on the predicted label (Figure 4). One hypothesis for the overall low similarity between models is that different models tend to give different predictions therefore they choose different features to support their decisions. However, we find that the similarity between models is not particularly high when they agree on the predicted label, and are sometimes

even lower than when they disagree. This is true for LIME in Yelp and for all methods in deception detection. In SST, the similarity when the models agree on the predicted label is generally greater than when they disagree. We show the comparison between SVM (ℓ_2) and LSTM here, and similar results hold for other combinations (see supplementary materials). This observation suggests that feature importance may not connect with the predicted labels: different models agree for different reasons and also disagree for different reasons.

The similarity between models and methods is generally negatively correlated with length but positively correlated with type-token ratio (Figure 5). Our results support **H2b**: Spearman correlation between length and similarity is mostly below 0, which indicates that the longer an instance is, the less similar the important features are. The negative correlation becomes stronger as k grows, indicating that length has a stronger effect on similarity when we consider more top features. However, this is not true in the case of LIME and SHAP where the correlation between length and similarity are occasionally above 0 and sometimes even




Figure 8: Distance of the part-of-speech tag distributions between important features and all words (background). Distance is generally smaller with post-hoc methods for all models, although some exceptions exist for LSTM with attention and BERT.

the declining relationship with k does not hold. Our result on type-token ratio is opposite to **H2c**: the greater the type-token ratio, the higher the similarity (see supplementary materials). We believe that the reason is that type-token ratio is strongly negatively correlated with length (the Spearman correlation for Yelp, SST and deception dataset is -0.92, -0.59 and -0.84 respectively). In other words, type-to-token ratio becomes redundant to length and fails to capture text complexity beyond length.

7 Distribution of Important Features

Finally, we examine the distribution of important features obtained from different approaches. These results may partly explain our previously observed low similarity in feature importance.

Important features show higher entropy using LSTM with attention and lower entropy with XGBoost (Figure 6). As expected from **H3a**, LSTM with attention (the pink lines) are usually at the top (similar results for BERT in the supplementary material). Such a high entropy can contribute to the low similarity between LSTM with attention and other models. However, as the order in similarity between SVM and XGBoost is less stable, entropy cannot be the sole cause.

Distribution of POS tags (Figure 7 and Figure 8). We further examine the linguistic properties of important words. Consistent with **H3b**, adjectives are more important in sentiment classification than in deception detection. On the contrary to our hypothesis, we found that pronouns do not always play an important role in deception detection. Notably, LSTM with attention puts a strong emphasis on nouns in deception detection. In all cases, determiners are under-represented among important words. With respect to the distance of part-

of-speech tag distributions between important features and all words (background), post-hoc methods tend to bring important words closer to the background words, which echoes the previous observation that post-hoc methods tend to increase the similarity between important words (Figure 8).

8 Concluding Discussion

In this work, we provide the first systematic characterization of feature importance obtained from different approaches. Our results show that different approaches can sometimes lead to very different important features, but there exist some consistent patterns between models and methods. For instance, deep learning models tend to generate diverse important features that are different from traditional models; post-hoc methods lead to more similar important features than built-in methods.

As important features are increasingly adopted for varying use cases (e.g., decision making vs. model debugging), we hope to encourage more work in understanding the space of important features, and how they should be used for different purposes. While we focus on consistent patterns across classification tasks, it is certainly interesting to investigate how properties related to tasks and data affect the findings. Another promising direction is to understand whether more concentrated important features (lower entropy) lead to better human performance in supporting decision making.

Acknowledgments

We thank Jim Martin, anonymous reviewers, and members of the NLP+CSS research group at CU Boulder for their insightful comments and discussions. This work was supported in part by NSF grant IIS-1849931.

References

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of ICLR*.

Tianqi Chen and Carlos Guestrin. 2016. XGBoost: A scalable tree boosting system. In *Proceedings of KDD*.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL*.

Finale Doshi-Velez and Been Kim. 2017. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*.

Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. 2019. A survey of methods for explaining black box models. *ACM computing surveys (CSUR)*, 51(5):93.

Matthew Honnibal and Ines Montani. 2017. spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing.

Benjamin D Horne, Dorit Nevo, John O’Donovan, Jin-Hee Cho, and Sibel Adali. 2019. Rating reliability and bias in news articles: Does ai assistance help everyone? In *Proceedings of ICWSM*.

Sarthak Jain and Byron C. Wallace. 2019. Attention is not explanation. In *Proceedings of NAACL*.

Thorsten Joachims. 1998. Text categorization with support vector machines: Learning with many relevant features. In *Proceedings of ECML*.

Been Kim, Rajiv Khanna, and Oluwasanmi O Koyejo. 2016. Examples are not enough, learn to criticize! criticism for interpretability. In *Proceedings of NeurIPS*.

Vivian Lai and Chenhao Tan. 2019. On human predictions with explanations and predictions of machine learning models: A case study on deception detection. In *Proceedings of FAT**.

Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2016. Visualizing and understanding neural models in nlp. In *Proceedings of NAACL*.

Zachary C Lipton. 2016. The mythos of model interpretability. *arXiv preprint arXiv:1606.03490*.

Tania Lombrozo. 2012. Explanation and abductive inference. *Oxford handbook of thinking and reasoning*, pages 260–276.

Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of NeurIPS*.

Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proceedings of EMNLP*.

Dmitry M Malioutov, Kush R Varshney, Amin Emad, and Sanjeeb Dash. 2017. Learning interpretable classification rules with boolean compressed sensing. In *Transparent Data Mining for Big and Small Data*, pages 95–121. Springer.

Dong Nguyen. 2018. Comparing automatic and human evaluation of local explanations for text classification. In *Proceedings of NAACL*.

Myle Ott, Claire Cardie, and Jeffrey T Hancock. 2013. Negative deceptive opinion spam. In *Proceedings of NAACL*.

Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T Hancock. 2011. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of ACL*.

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of KDD*.

Sofia Serrano and Noah A. Smith. 2019. Is Attention Interpretable? In *Proceedings of ACL*.

Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. Learning important features through propagating activation differences. In *Proceedings of ICML*.

Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of EMNLP*.

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 *Proceedings of NeurIPS*.

Jesse Vig. 2019. Deconstructing BERT: Distilling 6 Patterns from 100 Million Parameters. <https://towardsdatascience.com/\deconstructing-bert-distilling-6-patterns-from\100-million-parameters-b49113672f77>. [Online; accessed 27-Apr-2019].

Sarah Wiegreffe and Yuval Pinter. 2019. Attention is not not Explanation. In *Proceedings of EMNLP*.

Thomas Wolf. 2019. `huggingface/pytorch-pretrained-bert`. <https://github.com/huggingface/pytorch-pretrained-BERT>.

Yelp. 2019. Yelp dataset 2019. <https://www.yelp.com/dataset>. Accessed: 2019-01-01.