

A Song of (Dis)agreement: Evaluating the Evaluation of Explainable Artificial Intelligence in Natural Language Processing

Michael NEELY*[†] ^a, Stefan F. SCHOUTEN* ^{a,b}, Maurits BLEEKER ^a, and
Ana LUCIC ^a

^a *University of Amsterdam*

^b *Vrije Universiteit Amsterdam*

Abstract. There has been significant debate in the NLP community about whether or not attention weights can be used as an *explanation* – a mechanism for interpreting how important each input token is for a particular prediction. The validity of “attention as explanation” has so far been evaluated by computing the rank correlation between attention-based explanations and existing feature attribution explanations using LSTM-based models. In our work, we (i) compare the rank correlation between five more recent feature attribution methods and two attention-based methods, on two types of NLP tasks, and (ii) extend this analysis to also include transformer-based models. We find that attention-based explanations do not correlate strongly with any recent feature attribution methods, regardless of the model or task. Furthermore, we find that none of the tested explanations correlate strongly with one another for the transformer-based model, leading us to question the underlying assumption that we should measure the validity of attention-based explanations based on how well they correlate with existing feature attribution explanation methods. After conducting experiments on five datasets using two different models, we argue that the community should stop using rank correlation as an evaluation metric for attention-based explanations. We suggest that researchers and practitioners should instead test various explanation methods and employ a human-in-the-loop process to determine if the explanations align with human intuition for the particular use case at hand.

Keywords. Explainability, Interpretability, Natural Language Processing.

1. Introduction

As machine learning (ML) models are increasingly used in hybrid settings to make consequential decisions for humans, criteria for plausible and faithful explanations of their predictions remain speculative [1, 2]. Although there are many possible explanations for a model’s decision, only those faithful to both the model’s reasoning process and to hu-

* Equal contribution.

† Now at ASOS.com.

man stakeholders are desirable [2]. The rest are irrelevant in the best case and harmful in the worst, particularly in critical domains such as law [3], finance [4], and medicine [5].

Content moderation is a use case where explanations can help domain experts as part of a hybrid human-in-the-loop system. Consider an ML model that predicts whether or not a post on a social media contains misinformation: when a post is automatically removed by the model and the user who created it appeals its removal, content moderators need to read through the entire post to identify why it was flagged as containing misinformation [6]. Including explanations that identify which parts of the post are problematic can help content moderators decide if the model was correctly flagging the post.

ML practitioners frequently explain models by calculating each input’s contribution toward an individual prediction. Additivity — treating all contributions as independent and quantifiable — is a common simplifying assumption. In this work, we focus on such additive explanation methods and refer to them as *feature attribution methods*. We denote the contribution of each input toward the model’s decision as its *importance*. We say that two different feature attribution methods *agree* if there is a strong correlation between their computed rankings of input importance.

The attention mechanism [7] in natural language processing (NLP) is a popular, albeit less rigorously motivated, way of obtaining explanations. Because the mechanism produces context vectors from which decoders can soft-search for prediction-relevant information, the weights assigned to inputs intuitively serve as proxies for their overall contribution towards a decision. Weights are often visualized as heatmaps over sequences [e.g., 8, 9], which can be particularly persuasive when examples are (unintentionally) cherry-picked to fit a narrative. We define an attention-based explanation as a vector of attention weights that, similar to a feature attribution method, can be treated as a ranking of importance.

In their critique of attention-based explanations for NLP, Jain and Wallace [10] argue that faithful attention-based explanations must be highly *agreeable*.¹ That is, their generated rankings of input importance must correlate with existing feature attribution methods. Following Jain and Wallace [10] and their claim that “attention is not explanation”, several recent papers have presented an increased agreement with a small set of feature attribution methods as evidence for their proposed method’s ability to improve the faithfulness of the attention mechanism. For example, Mohankumar et al. [12] show that minimizing hidden state concity in a BiLSTM improves the Pearson correlation of attention weights with Integrated Gradients [13] attributions. As the popularity of *agreement as evaluation* grows [14, 15], we believe it is important to investigate diagnostic capacity of agreement as a metric by examining (i) more recent feature attribution methods, and (ii) more complex transformer-based models.

Under the paradigm of *agreement as evaluation*, new explanation methods (i.e., attention-based) are compared to established explanation method(s) (i.e., feature attributions). However, can any one explanation method act as the standard against which other explanation methods are evaluated? Explanations are task-, model-, and context-specific [16], and the performance of explanation methods depends on the particular diagnostic tests considered [17, 18].

In this work, our main research question is: *How well do attention-based explanations correlate with recent feature attribution methods for NLP tasks?* Specifically, we want to investigate:

¹Ethayarajh and Jurafsky [11] use the term *consistent*.

- **RQ1:** Does the correlation depend on the model architecture (transformer- vs. LSTM-based)?
- **RQ2:** Does the correlation depend on the nature of the classification task (single- and pair-sequence)?

We investigate the following feature attribution methods: LIME [19], Integrated Gradients [13], DeepLIFT [20], and two versions of SHAP [21]: Grad-SHAP (based on Integrated Gradients) and Deep-SHAP (based on DeepLIFT). We observe low agreement between attention-based explanations and feature attribution methods, across both models and both tasks. We also observe low agreement across all explanation methods for the transformer-based model and for pair-sequence tasks. We use this empirical evidence, along with our theoretical objections, to argue that practitioners should refrain from evaluating attention-based explanations based on their agreement with feature attribution methods.

2. Related Work

Jain and Wallace [10] introduced the *agreement as evaluation* paradigm by comparing attention-based explanations with simple feature attribution methods using a bidirectional LSTM [22] on single- and pair-sequence classification tasks. They conclude that “attention is not explanation” due to the weak correlation between the rankings of input token importance obtained from attention weights and those obtained from two elementary feature attribution methods: (i) input \times gradient [23, 24], and (ii) leave-one-out [25]. In our work, we test the generalizability of agreement as an evaluation metric by (i) testing on a more comprehensive set of feature attribution methods (see Section 3.1), and (ii) testing on a transformer-based [26] model on the same types of tasks (see Section 4.2).

The influential critique by Jain and Wallace [10] has sparked an ongoing debate about whether or not attention is explanation [27, 28, 29]. More recently, Bastings and Filippova [30] have questioned the notion of “attention as explanation” as a whole, and suggest that in order to explain ML model predictions, the community should rely on methods that are explicitly created for this purpose (i.e., feature attributions), instead of seeking explanations from attention mechanisms. In our work, we do not aim to take a position on the “is attention explanation” debate, but rather investigate the hypothesis that in order for attention mechanisms to be considered “explanations”, they must correlate with existing feature attribution methods. This is an underlying assumption of not only the work by Jain and Wallace [10], but also that of Meister et al. [14], who show that inducing sparsity in the attention distribution decreases agreement with feature attribution methods, and Abnar and Zuidema [15], who demonstrate their *attention-flow* algorithm improves the correlation with attributions based on feature ablation.

Atanasova et al. [31] introduce a series of diagnostic tests to evaluate feature attribution methods for text classification. They show that the performance of feature attribution methods, measured by using these diagnostic tests, largely depends on the model and task considered, but note that gradient-based methods tend to perform the best. Similar to their work, we also compare and evaluate feature attribution methods, but only to investigate the suitability of *agreement as evaluation*, not to determine a winning explanation method given several diagnostic tests. In contrast, we evaluate five feature at-

tribution methods on agreement as evaluation based on only the rank correlation using Kendall's- τ , following Jain and Wallace [10].

Prasad et al. [32] define three alignment metrics that quantify how well human-annotated natural language explanations align with the explanations generated by the Integrated Gradients method [13]. They find that the BERT [33] model has the highest alignment with human-annotated explanations. However, unlike our work, this work focuses exclusively on using transformer-based models and Integrated Gradients, and does not question the *agreement as evaluation* paradigm for feature attribution methods.

Ding and Koehn [34] introduce a human-annotated benchmark to evaluate feature attribution methods for NLP models. They test two attribution methods on three types of NLP models and find that explanations from feature attribution methods are sensitive to changes in model configuration. Similarly, we test five attribution methods on two types of NLP models, but our focus is on investigating the *agreement as evaluation* paradigm, whereas Ding and Koehn [34] focus on investigating the correlation between feature attribution methods and a human-annotated benchmark. We are unable to incorporate the benchmark proposed by Ding and Koehn [34] in our work because the human annotations are not present for every single token, and therefore we cannot turn them into a ranking.

This work builds off our prior ICML workshop paper [35], which several other papers have extended. Feldhus et al. [36] introduce a software package to analyze instance-wise explanations for popular NLP models and tasks and, in doing so, partially reproduce one of our experiments. Krishna et al. [37] formally define and highlight the importance of the "disagreement problem" between feature attribution methods, which they find is a constant frustration for ML practitioners. They introduce metrics to capture the disagreement between top-k features and "features of interest" (e.g., those selected by an end-user) and find considerable disagreement between feature attribution methods for tabular, text, and image data on both real-world and research datasets.

3. Explainability in NLP

In this work we investigate two types of explanations for NLP tasks: (i) those from recent feature attribution methods in the explainable AI (XAI) literature and (ii) those based on attention scores. We evaluate on transformer- and LSTM-based models (**RQ1**), for both single- and pair-sequence tasks (**RQ2**). Following Jain and Wallace [10], we define an *explanation* of an input token sequence as a vector comprised of an importance score for each token. This vector can be used to rank the tokens by their importance score.

3.1. Explanations from Feature Attributions

In our experiments, we focus on five recent feature attribution methods: LIME [19], Integrated Gradients [13], DeepLIFT [20], and two versions of SHAP [21]: Grad-SHAP and Deep-SHAP:

- LIME [19] produces locally faithful explanations by learning an interpretable (e.g., linear) model from samples weighted by their proximity to the original instance.
- Integrated Gradients [13] calculates input feature attributions by accumulating the gradients obtained from the model along the straight-line path from a baseline to the original instance.

- DeepLIFT [20] also produces input feature attributions using the gradients, but it assigns scores based on the difference between a reference activation and the activation of the original instance. This allows the calculated contributions to remain non-zero even when the gradients are zero.
- SHAP [21] identifies a unique solution for the contribution of each input toward the prediction. Since this is computationally expensive, Lundberg and Lee [21] propose approximations based on existing methods:
 - Grad-SHAP (based on Integrated Gradients).
 - Deep-SHAP (based on DeepLIFT).

3.2. Explanations from Attention Mechanisms

Given an input sequence of tokens $\mathbf{S} = (t_1, \dots, t_n)$, we define an *attention-based explanation* as an assignment of attention weights $\alpha \in \mathbb{R}^n$ to the tokens in \mathbf{S} [10]. Since the dimensionality of α is architecture-dependent, it may be necessary to filter or aggregate the weights. In our experiments, this is only relevant for the self-attention mechanism in the transformer-based model we consider below (see Section 4.2).

Previous analyses at the attention head level implicitly assume that contextual word embeddings remain tied to their corresponding tokens across self-attention layers [38, 39]. This assumption may not hold in transformer-based models since information mixes across layers [40]. Therefore, we use the *attention rollout* [15] method — which assumes the identities of tokens are linearly combined through the self-attention layers based exclusively on attention weights — to calculate post-hoc, token-level importance scores.² Following Abnar and Zuidema [15], we use the scores calculated for the last layer’s [CLS] token, resulting in a final vector $\alpha \in \mathbb{R}^n$ at the time of evaluation.

Recurrent models similarly suffer from issues of identifiability. In LSTM-based models, attention is computed over hidden representations across timesteps, which does not provide faithful token-level importance scores. Approaches that trace explanations back to individual timesteps [41] or input tokens [42] are only just emerging. Therefore, we analyze the raw attention weights for the LSTM-based model we consider below (see Section 4.2).

3.3. Agreement between Explanations

Following Jain and Wallace [10], we measure *agreement* between the explanation methods as the mean Kendall- τ correlation [43] between the ranked importance scores of all input tokens, across all examples.³ The Kendall- τ correlation is a widely used metric for comparing ranked lists; it measures the correlation between two ranked lists based on discordant pairs between the lists. Two items are considered *discordant* if they are ranked differently on the two lists. The Kendall- τ correlation can take on values in the $[-1, 1]$ interval, where negative values imply the rankings are negatively correlated while positive values imply the rankings are positively correlated [43].

²We also experimented with *attention flow* [15] with similar results (see our public repository).

³We also calculated Spearman and Pearson correlations and obtained similar results.

4. Experimental Setup

4.1. Datasets

We evaluate on two types of NLP classification tasks: (i) single-sequence, and (ii) pair-sequence. For the single-sequence task, we perform binary sentiment classification on the Stanford Sentiment Treebank (SST-2) [44] and the IMDb Large Movie Reviews Corpus (IMDb) [45]. We use identical splits and pre-processing as Jain and Wallace [10], but also remove sequences longer than 240 tokens for faster attribution calculation. This leaves us with roughly 78% of the original instances in the IMDb dataset.

- The SST-2 dataset consists of single sentences extracted from movie reviews and is used for binary sentiment analysis: predicting whether the review is positive or negative.
- The IMDb dataset consists of movie reviews as well, but contains longer sequences compared to SST-2. It is also used for binary sentiment analysis.

For the pair-sequence task, we examine natural language inference and understanding with the Multi Natural Language Inference (MNLI) corpus [46], the Stanford Natural Language Inference (SNLI) corpus [47], and the Quora Question Pairs dataset [48].

- The MNLI dataset contains sentence pairs for a textual entailment task: given a pair of sentences, we want to predict whether or not one sentence implies the other. Since MNLI has no publicly available test set, we use the English subset of the XNLI [49] test set.
- The SNLI dataset [47] contains pairs of sentences and is used for the textual entailment task, similar to MNLI.
- The Quora Question Pairs dataset [48] contains pairs of questions from the Quora website, where the task is to classify if the two questions in a pair are duplicates or not. We use a custom split (80/10/10) for the Quora dataset, removing pairs with a combined count of 200 or more tokens (leaving 99.99% of the original instances).

4.2. Models

We test two types of NLP models: (i) transformer-based, and (ii) LSTM-based. The transformer-based model [26] relies on a self-attention mechanism to calculate representations for tokens. For each layer in the transformer network, the representation of each token is updated by computing a weighted sum over all tokens represented in the entire sequence, together with a non-linear feedforward transformation. The weight value of each token is determined using self-attention, which computes a similarity score for each pair of tokens in the sequence. For every token in every layer of the network, an attention layer is used. As a result, it is non-trivial to aggregate all the importance values per layer into one interpretable importance score per token.

In our work, we follow Abnar and Zuidema [15] by using the attention rollout scores for [CLS] tokens. For the transformer-based model, we fine-tune the lighter, pre-trained DistilBERT variant [50] instead of the full BERT model [33] to reduce the computational overhead and ecological footprint. For classification, we add a linear layer on top of the pooled output. We concatenate pair-sequences with a [SEP] token.

Unlike a transformer-based model, an LSTM-based model processes input tokens in sequential order. For each token, $t \in \{1, \dots, n\}$, in the sequence, the global representation of the sequence, \mathbf{h}_t , is updated using a non-linear transformation with an embedded representation of (i) the current token in the sequence, t_n , and (ii) the previous global representation of the sequence, \mathbf{h}_{t-1} , as input. A bidirectional LSTM (Bi-LSTM) uses two stacked LSTMs to process tokens in both directions of the sequence.

For the Bi-LSTM, we use the same single-layered bidirectional encoder with the query-less additive (*tanh*) attention and linear feedforward decoder as in the work of Jain and Wallace [10]. As a result, the attention weight for each token in the sequence is solely based on its representation \mathbf{h}_t , and not on a query item. In pair-sequence tasks, we embed, encode, and induce attention over each sequence separately. The decoder predicts the label from the concatenation of (i) both context vectors c_1 and c_2 , (ii) their absolute difference $|c_1 - c_2|$, and (iii) their element-wise product $c_1 \odot c_2$.

4.3. Training the models

We train three independently-seeded instances of each of the models described in Section 4.2 using the AllenNLP framework [51], each for a maximum of 40 epochs. We use a patience value of 5 epochs for early stopping. For DistilBERT, we fine-tune the standard “base-uncased” weights available in the HuggingFace library [52] with the AdamW [53] optimizer. For the BiLSTM, we follow Jain and Wallace [10] and select a 128-dimensional encoder hidden state with a 300-dimensional embedding layer. We tune pre-trained FastText embeddings [54] and optimize with the AMSGrad variant [55] of Adam [56].

4.4. Explaining the models

We leverage the Captum⁴ implementations of LIME, Integrated Gradients, DeepLIFT, Grad-SHAP, and Deep-SHAP, and use the padding token as a baseline where applicable. For LIME, we mask tokens as features and use 1000 samples to train the interpretable models. We apply our feature attribution methods to the predictions of each independently-seeded model for 500 instances randomly sampled from each test set.

Our code is available in our public repository⁵, which also contains more information on reproducing our experiments.

5. Results

In this section, we answer our main research question: *How well do attention-based explanations correlate with recent feature attribution methods for NLP tasks?* In general, we find that attention-based explanations do not correlate strongly with feature attribution methods, with some exceptions (see Section 5.2 and 5.3).

⁴github.com/pytorch/captum

⁵github.com/sfschouten/court-of-xai

5.1. RQ1: Does the correlation depend on the model architecture?

Tables 1 and 2 display the average⁶ Kendall- τ correlations between the explanation methods for the DistilBERT and BiLSTM models, respectively. The same tables but including standard deviations are available in our public repository. A stronger correlation (i.e., agreement) is indicated by a darker blue colour in the table cell. In general, we see that the agreement between explanation methods is substantially lower for the DistilBERT model than for the BiLSTM model.

Table 1. Mean Kendall- τ between the tested explanation methods for the DistilBERT model. A darker color indicates a stronger correlation between the compared explanation methods. Attn Roll refers to Attention Rollout.

		Attn Roll	LIME	Int-Grad	DeepLIFT	Grad-SHAP	Deep-SHAP
Attn Roll	IMDb	1	.1259	.1818	.2516	.1432	.2303
	SST-2	1	.1359	.0511	.1328	.0737	.1291
	MNLI	1	.2678	.1891	.2432	.1905	.2067
	Quora	1	.1622	.0574	.2267	.0518	.2257
	SNLI	1	.1434	.1645	.2214	.1600	.1796
LIME	IMDb		1	.1050	.0696	.0929	.0655
	SST-2		1	.2861	.0618	.2414	.0499
	MNLI		1	.1794	.1526	.1592	.1205
	Quora		1	.1407	.0032	.1144	.0095
	SNLI		1	.1529	.0925	.1104	.0593
Int-Grad	IMDb			1	.1433	.5495	.1246
	SST-2			1	.0498	.4987	.0381
	MNLI			1	.2153	.4780	.1708
	Quora			1	.0625	.4674	.0529
	SNLI			1	.0955	.3932	.0700
DeepLIFT	IMDb				1	.1306	.4830
	SST-2				1	.0522	.4514
	MNLI				1	.2324	.4985
	Quora				1	.0637	.5951
	SNLI				1	.1181	.5554
Grad-SHAP	IMDb					1	.1093
	SST-2					1	.0419
	MNLI					1	.1752
	Quora					1	.0535
	SNLI					1	.0851
Deep-SHAP	IMDb						1
	SST-2						1
	MNLI						1
	Quora						1
	SNLI						1

For the DistilBERT model, we observe a weak correlation across all explanations – almost none of the explanation methods we test, whether they are attention-based or feature attributions, correlate strongly with one another. There are two exceptions: (i) Integrated Gradients moderately correlates with Grad-SHAP, and (ii) DeepLIFT moderately correlates with Deep-SHAP. However, this is unsurprising since the implementation of Grad-SHAP is based on Integrated Gradients, and the implementation of Deep-SHAP is based on DeepLIFT. In contrast, it is surprising to see the lack of correlation between

⁶Across the 3 model instances, randomly selecting 500 instances from the test set using the training seed.

Grad-SHAP and Deep-SHAP for the DistilBERT model, given that they are different implementations of the same algorithm, SHAP [21].

Table 2. Mean Kendall- τ between the tested explanation methods for the BiLSTM. A darker color indicates a stronger correlation between the compared explanation methods.

		Attn Weights	LIME	Int-Grad	DeepLIFT	Grad-SHAP	Deep-SHAP
Attn Weights	IMDb	1	.2014	.2188	.2494	.2209	.2309
	SST-2	1	.1326	.1093	.1372	.1101	.1400
	MNLI	1	.1958	.2523	.2549	.2473	.2370
	Quora	1	.0363	.0143	.0894	.0182	.1017
	SNLI	1	.2198	.2566	.3158	.2517	.2938
LIME	IMDb		1	.6538	.5854	.6486	.5584
	SST-2		1	.4968	.4734	.4962	.4422
	MNLI		1	.3281	.2444	.3187	.2269
	Quora		1	.2099	.1900	.2037	.1670
	SNLI		1	.2673	.1676	.2481	.1566
Int-Grad	IMDb			1	.7331	.9409	.6994
	SST-2			1	.8683	.9707	.8063
	MNLI			1	.4984	.8138	.4021
	Quora			1	.2906	.7420	.2290
	SNLI			1	.2461	.6535	.2165
DeepLIFT	IMDb				1	.7378	.8593
	SST-2				1	.8682	.8729
	MNLI				1	.4987	.6208
	Quora				1	.3158	.6179
	SNLI				1	.2557	.5791
Grad-SHAP	IMDb					1	.7021
	SST-2					1	.8056
	MNLI					1	.4015
	Quora					1	.2433
	SNLI					1	.2219
Deep-SHAP	IMDb						1
	SST-2						1
	MNLI						1
	Quora						1
	SNLI						1

For the BiLSTM model, we observe a weak correlation between attention-based explanations and feature attribution explanations (i.e., the first row of Table 2). Similar to the DistilBERT model, we see strong correlations for the methods that have similar underlying implementations. We also see some strong correlation between feature attribution methods, especially for single-sequence tasks (see Section 5.2).

Overall, we conclude that the correlation between explanation methods depends on the model architecture, which answers **RQ1**. In general, the correlation is weaker for the DistilBERT model than for the BiLSTM. However, the overall agreement between attention-based explanations and feature attribution explanations is weak, regardless of the model architecture.

5.2. RQ2: Does the correlation depend on the nature of the classification task?

To investigate **RQ2**, we examine five different datasets corresponding to two different tasks: IMDb and SST-2 are single-sequence tasks, while MNLI, Quora, and SNLI are pair-sequence tasks and show the results in Tables 1 and 2.

For the BiLSTM model, the results are clear: there is a substantially stronger correlation among explanations for single-sequence tasks than for pair-sequence tasks. For

the DistilBERT model however, we do not see such a clear distinction – there is, at best, a weak correlation across all explanation methods, regardless of the nature of the task.

Regarding **RQ2**, we conclude that explanation methods agree with each other more on single-sequence tasks than pair-sequence tasks. This relationship holds especially for the BiLSTM model compared to the DistilBERT model.

5.3. Explanations that strongly correlate

Given that some of the feature attribution methods are inherently related, we would expect to see stronger degrees of correlation between them: (i) Grad-SHAP and Deep-SHAP are two different versions of the SHAP explanation method [21], (ii) Grad-SHAP relies on the Integrated Gradients implementation to compute the token importances, and (iii) Deep-SHAP relies on DeepLIFT. We see some strong correlations for (ii) and (iii) between these methods for both models. These correlations are especially strong for the single-sequence tasks on the BiLSTM. However, we see remarkably weak correlations for (i) on DistilBERT, and for the BiLSTM on the pair-sequence tasks.

6. Discussion

Based on the results in Section 5, we argue that rank correlation with existing feature attribution methods is not an appropriate measure for evaluating attention-based explanations. In this section, we detail three main reasons for this: (i) the general lack of correlation across all explanation methods, especially for transformer-based models, (ii) the fact that similar explanations do not always result in correlated rankings, and (iii) the lack of justification for the existence of one “ideal” explanation, which is a fundamental assumption of the *agreement as evaluation* paradigm.

6.1. Lack of correlation between explanation methods.

In Section 5, we have shown that there is a low degree of correlation between explanation methods, especially for the transformer-based model.⁷ Similar conclusions are observed in the work of [36] and [37]. This makes it challenging to justify the expectation that in order for attention-based explanations to be valid, they should correlate with existing feature attribution methods. If none of the existing feature attribution methods correlate with one another (as is the case for the transformer-based model), should we expect attention-based explanations to correlate with them? Therefore, we argue against the use of rank correlation as an “off-the-shelf” tool to evaluate attention-based explanations.

A common critique of using Kendall- τ is that expecting the full rankings of token importance to correlate is unrealistic. In their original paper, Jain and Wallace [10] theorized that top- k token comparison would reduce noise, improve correlation, and more closely align with the end user’s interest in the most salient features. However, formulating a systematic method for calculating k across various models, tasks, and datasets is difficult in practice and yields mixed results. Treviso and Martins [57] show a dynamic

⁷Although there are some examples of stronger correlations, these are either on a very specific combination of task and model (i.e., single sequence tasks on the BiLSTM), or due to similarities in the underlying implementation of the explanation methods.

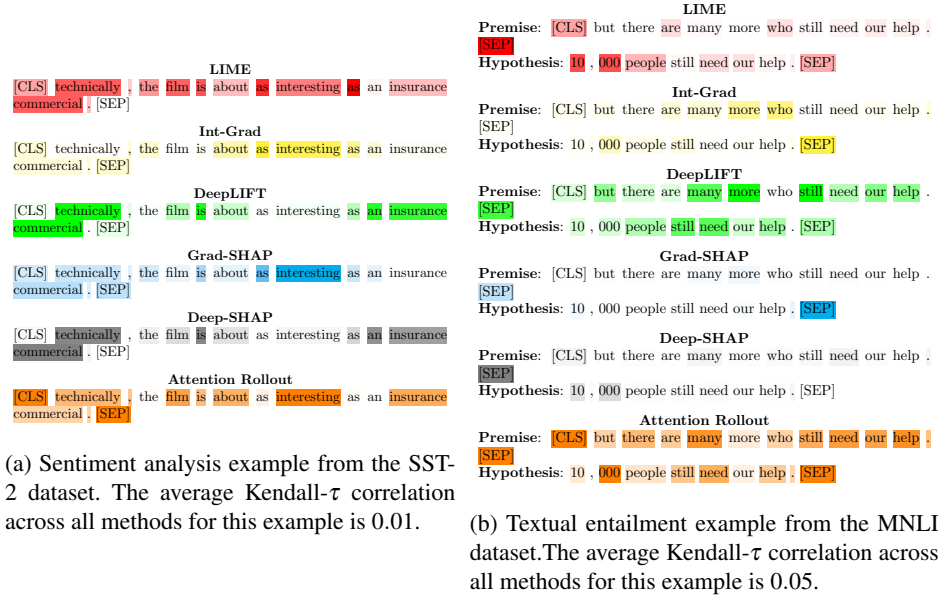


Figure 1. Examples of explanations for the transformer-based model. The brighter the color, the higher the attribution value

selection of k with a sparse attention mechanism more effectively conveys justifications of decisions than fixed values of k (e.g., 5 or 10 tokens). In contrast, Krishna et al. [37] show a widespread disagreement of feature attribution methods across domains regardless of whether k is fixed or dynamic (e.g., the top $X\%$ of features or top $X\%$ of tokens based on average sentence length).

6.2. Similar explanations do not imply strongly correlated rankings

In order to investigate what weak correlation looks like in practice, we apply a heatmap of importance scores for each of the tested explanation methods. The darker the color, the more important the token is w.r.t. the model prediction. Figure 1a shows a sentiment analysis example using the transformer-based model on the SST-2 dataset. Upon first glance, the visualizations in Figure 1a look relatively similar: the methods are all highlighting roughly the same words – almost all methods (except for Int-Grad) highlight the word ‘technically’ as important. However, the average Kendall- τ correlation across all methods for this example is very weak – only 0.01. This highlights the problem of evaluating explanations by measuring the correlation between two ranked lists of importance scores: two explanation methods can roughly indicate the same token(s) as important, but when similar tokens do not end up in similar positions in the two rankings, the overall agreement between the two methods can still be low.

6.3. Is there one ideal explanation?

The *agreement as evaluation* paradigm implicitly assumes the existence of a single “ideal” explanation (i.e. ranking of tokens) which all methods must uncover, and new feature attribution methods are evaluated based on how strongly they correlate with the

“ideal” explanation. However, it is unclear whether this assumption holds. For instance, input token importance rankings may only capture a narrow slice of the model’s behavior, such that many plausible rankings exist.

Since many tasks may be too complex for humans to judge token-level importance, it can be unclear how to choose the “ideal” ranking. While a handful of highly polar tokens are generally indicative of the label in binary sentiment classification [58], annotators may be unsure how to rank the other tokens. The difficulty increases in the pair-sequence setting – if two words indicate a contradiction, which is more important? As a result, when agreement is measured in the presence of multiple faithful and plausible rankings, feature attribution methods may look deceptively problematic, even if they are not.

Figure 1b shows an example of how different explanation methods can highlight different tokens as being important for the prediction, even though the underlying model is the same. Figure 1b shows a textual entailment example from the MNLI dataset, where the task is to predict whether or not the second sentence entails the first. The first sentence comes after the [CLS] token: “*But there are many more who still need our help.*”. The second sentence comes after the first [SEP] token: “*10,000 people still need our help*”. Given that we see different explanations for the same example using the same model, it is unclear how to identify which of these explanations is the “ideal” one.

Instead of evaluating potential explanations based on how well they correlate with one another, we suggest that researchers and practitioners should test various explanation methods and employ a human-in-the-loop process to determine if the explanations align with human intuition for the particular use case at hand. This can be used in combination with metrics such as those suggested by Atanasova et al. [31].

7. Conclusion

In this work, we investigate how strongly attention-based explanations correlate with recent feature attribution methods for NLP tasks. We compare attention-based explanations to five recent feature attribution methods, using both transformer- and LSTM-based models, on both single- and pair-sequence tasks. Overall, we observe a low degree of correlation between the attention-based explanations and the feature attribution explanations. For the transformer-based model, we find a weak correlation across all explanations for both task types. For the BiLSTM, we observe some strong correlations between explanations, but only for simple single-sequence tasks.

Through our experiments, we discover (i) a general lack of correlation between explanation methods, especially for more complex settings (i.e., transformer-based model and pair-sequence tasks) which is corroborated by additional recent research on text, tabular, and image data [36, 37], (ii) that similar explanations do not always result in correlated rankings, and (iii) the existence of a single “ideal” explanation is questionable, which is a fundamental assumption of the *agreement as evaluation* paradigm. Without an external ground-truth explanation, all that rank correlation tells us is whether or not two rankings are similar. For this reason, we recommend practitioners stop using *agreement as evaluation* for attention-based explanations.

In future work, we plan to approach this problem from first principles by formulating toy data to guarantee a single, correct top- k or full ranking for each instance to see if XAI methods are consistently capable of recovering the ideal explanation under this setting.

Acknowledgements

We thank the anonymous reviewers, Bilal Alsallakh and Maarten de Rijke for their helpful discussions and suggestions. This research was supported by Ahold Delhaize and the Netherlands Organisation for Scientific Research under project nr. 652.001.003 and the Nationale Politie. All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.

References

- [1] Zachary C. Lipton. The mythos of model interpretability: In machine learning, the concept of interpretability is both important and slippery. *Queue*, 16(3):31?57, June 2018. ISSN 1542-7730. doi: 10.1145/3236386.3241340. URL <https://doi.org/10.1145/3236386.3241340>.
- [2] Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4198–4205, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.386. URL <https://aclanthology.org/2020.acl-main.386>.
- [3] Danielle Leah Kehl and Samuel Ari Kessler. Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing. In *Responsive Communities Initiative*. Berkman Klein Center for Internet & Society, 2017.
- [4] Rory Mc Grath, Luca Costabello, Chan Le Van, Paul Sweeney, Farbod Kamiab, Zhao Shen, and Freddy Lecue. Interpretable Credit Application Predictions With Counterfactual Explanations. In *HAL*, December 2018. URL <https://hal.inria.fr/hal-01934915>.
- [5] Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '15*, pages 1721–1730, New York, NY, USA, 2015. Association for Computing Machinery. ISBN 9781450336642. doi: 10.1145/2783258.2788613. URL <https://doi.org/10.1145/2783258.2788613>.
- [6] Alon Halevy, Cristian Canton-Ferrer, Hao Ma, Umut Ozertem, Patrick Pantel, Marzieh Saeidi, Fabrizio Silvestri, and Ves Stoyanov. Preserving integrity in online social networks. *Communications of the ACM*, 65(2):92–98, January 2022. ISSN 0001-0782. doi: 10.1145/3462671. URL <https://doi.org/10.1145/3462671>.
- [7] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1409.0473>.
- [8] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In Francis Bach and David Blei, editors, *Proceedings of the 32nd International Conference on Machine Learning*, vol-

- ume 37 of *Proceedings of Machine Learning Research*, pages 2048–2057, Lille, France, 07–09 Jul 2015. PMLR. URL <https://proceedings.mlr.press/v37/xuc15.html>.
- [9] Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 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*, pages 681–691, San Diego, California, June 2016. Association for Computational Linguistics. doi: 10.18653/v1/N16-1082. URL <https://aclanthology.org/N16-1082>.
- [10] Sarthak Jain and Byron C. Wallace. Attention is not Explanation. 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 3543–3556, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1357. URL <https://aclanthology.org/N19-1357>.
- [11] Kawin Ethayarajh and Dan Jurafsky. Attention flows are shapley value explanations. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 49–54, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-short.8. URL <https://aclanthology.org/2021.acl-short.8>.
- [12] Akash Kumar Mohankumar, Preksha Nema, Sharan Narasimhan, Mitesh M. Khapra, Balaji Vasanth Srinivasan, and Balaraman Ravindran. Towards transparent and explainable attention models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4206–4216, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.387. URL <https://aclanthology.org/2020.acl-main.387>.
- [13] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML’17, pages 3319–3328. JMLR.org, 2017. URL <http://proceedings.mlr.press/v70/sundararajan17a.html>.
- [14] Clara Meister, Stefan Lazov, Isabelle Augenstein, and Ryan Cotterell. Is sparse attention more interpretable? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 122–129, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-short.17. URL <https://aclanthology.org/2021.acl-short.17>.
- [15] Samira Abnar and Willem Zuidema. Quantifying attention flow in transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.385. URL <https://aclanthology.org/2020.acl-main.385>.
- [16] Finale Doshi-Velez and Been Kim. Towards A Rigorous Science of Interpretable Machine Learning. *arXiv*, 2017. URL <https://arxiv.org/abs/1702.08608>.
- [17] Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. ERASER: A benchmark to evaluate rationalized NLP models. In *Proceedings of the 58th Annual Meeting of the Associa-*

- tion for Computational Linguistics, pages 4443–4458, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.408. URL <https://aclanthology.org/2020.acl-main.408>.
- [18] Marko Robnik-Šikonja and Marko Bohanec. *Perturbation-Based Explanations of Prediction Models*, pages 159–175. Springer International Publishing, Cham, 2018. ISBN 978-3-319-90403-0. doi: 10.1007/978-3-319-90403-0_9. URL https://doi.org/10.1007/978-3-319-90403-0_9.
- [19] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 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*, KDD '16, pages 1135–1144, New York, NY, USA, 2016. Association for Computing Machinery. ISBN 9781450342322. doi: 10.1145/2939672.2939778. URL <https://doi.org/10.1145/2939672.2939778>.
- [20] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ICML'17, pages 3145–3153. JMLR.org, 2017. URL <http://proceedings.mlr.press/v70/shrikumar17a.html>.
- [21] Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4765–4774, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>.
- [22] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- [23] Pieter-Jan Kindermans, Kristof Schütt, Klaus-Robert Müller, and Sven Dähne. Investigating the influence of noise and distractors on the interpretation of neural networks. *CoRR*, abs/1611.07270, 2016. URL <http://arxiv.org/abs/1611.07270>.
- [24] Yotam Hechtlinger. Interpretation of prediction models using the input gradient. *CoRR*, abs/1611.07634, 2016. URL <http://arxiv.org/abs/1611.07634>.
- [25] Jiwei Li, Will Monroe, and Dan Jurafsky. Understanding neural networks through representation erasure. *CoRR*, abs/1612.08220, 2016. URL <http://arxiv.org/abs/1612.08220>.
- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Isabelle Guyon, Ulrike von Luxburg, Samy Bengio, Hanna M. Wallach, Rob Fergus, S. V. N. Vishwanathan, and Roman Garnett, editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html>.

- [27] Sarah Wiegrefe and Yuval Pinter. Attention is not not explanation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China, November 2019. Association for Computational Linguistics. doi: 10.18653/v1/D19-1002. URL <https://aclanthology.org/D19-1002>.
- [28] Sofia Serrano and Noah A. Smith. Is attention interpretable? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1282. URL <https://aclanthology.org/P19-1282>.
- [29] Christopher Grimsley, Elijah Mayfield, and Julia R.S. Bursten. Why attention is not explanation: Surgical intervention and causal reasoning about neural models. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1780–1790, Marseille, France, May 2020. European Language Resources Association. ISBN 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.220>.
- [30] Jasmijn Bastings and Katja Filippova. The elephant in the interpretability room: Why use attention as explanation when we have saliency methods? In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 149–155, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.blackboxnlp-1.14. URL <https://aclanthology.org/2020.blackboxnlp-1.14>.
- [31] Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. A diagnostic study of explainability techniques for text classification. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3256–3274, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.263. URL <https://aclanthology.org/2020.emnlp-main.263>.
- [32] Grusha Prasad, Yixin Nie, Mohit Bansal, Robin Jia, Douwe Kiela, and Adina Williams. To what extent do human explanations of model behavior align with actual model behavior? In *Proceedings of the Fourth BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 1–14, Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.blackboxnlp-1.1. URL <https://aclanthology.org/2021.blackboxnlp-1.1>.
- [33] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1423. URL <https://aclanthology.org/N19-1423>.
- [34] Shuoyang Ding and Philipp Koehn. Evaluating saliency methods for neural language models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5034–5052, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.399. URL <https://aclanthology.org/2021.naacl-main.399>.

- [35] Michael Neely, Stefan F. Schouten, Maurits J. R. Bleeker, and Ana Lucic. Order in the court: Explainable ai methods prone to disagreement. In *ICML 2021 Workshop on Theoretic Foundation, Criticism, and Application Trend of Explainable AI*. International Conference on Machine Learning, July 2021. doi: <https://doi.org/10.48550/arXiv.2107.08821>. URL <https://arxiv.org/abs/2105.03287>.
- [36] Nils Feldhus, Robert Schwarzenberg, and Sebastian Möller. Thermostat: A large collection of NLP model explanations and analysis tools. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 87–95, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-demo.11. URL <https://aclanthology.org/2021.emnlp-demo.11>.
- [37] Satyapriya Krishna, Tessa Han, Alex Gu, Javin Pombra, Shahin Jabbari, Steven Wu, and Himabindu Lakkaraju. The disagreement problem in explainable machine learning: A practitioner’s perspective, 2022. URL <https://arxiv.org/abs/2202.01602>.
- [38] Joris Baan, Maartje ter Hoeve, Marlies van der Wees, Anne Schuth, and Maarten de Rijke. Understanding multi-head attention in abstractive summarization. *CoRR*, abs/1911.03898, 2019. URL <http://arxiv.org/abs/1911.03898>.
- [39] Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 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, August 2019. Association for Computational Linguistics. doi: 10.18653/v1/W19-4828. URL <https://aclanthology.org/W19-4828>.
- [40] Gino Brunner, Yang Liu, Damian Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. On identifiability in transformers. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=BJg1f6EFDB>.
- [41] João Bento, Pedro Saleiro, André F. Cruz, Mário A.T. Figueiredo, and Pedro Bizarro. TimeSHAP: Explaining Recurrent Models through Sequence Perturbations. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, KDD ’21*, pages 2565–2573, New York, NY, USA, August 2021. Association for Computing Machinery. ISBN 978-1-4503-8332-5. doi: 10.1145/3447548.3467166. URL <https://doi.org/10.1145/3447548.3467166>.
- [42] Martin Tutek and Jan Snajder. Staying true to your word: (how) can attention become explanation? In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 131–142, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.repl4nlp-1.17. URL <https://aclanthology.org/2020.repl4nlp-1.17>.
- [43] M. G. KENDALL. A new measure of rank correlation. *Biometrika*, 30(1-2):81–93, 06 1938. ISSN 0006-3444. doi: 10.1093/biomet/30.1-2.81. URL <https://doi.org/10.1093/biomet/30.1-2.81>.
- [44] Richard Socher, John Bauer, Christopher D. Manning, and Andrew Y. Ng. Parsing with compositional vector grammars. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages

- 455–465, Sofia, Bulgaria, August 2013. Association for Computational Linguistics. URL <https://aclanthology.org/P13-1045>.
- [45] Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <https://aclanthology.org/P11-1015>.
- [46] Adina Williams, Tiago Pimentel, Hagen Blix, Arya D. McCarthy, Eleanor Chodroff, and Ryan Cotterell. Predicting declension class from form and meaning. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6682–6695, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.597. URL <https://aclanthology.org/2020.acl-main.597>.
- [47] Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1075. URL <https://aclanthology.org/D15-1075>.
- [48] Lakshay Sharma, Laura Graesser, Nikita Nangia, and Utku Evci. Natural language understanding with the quora question pairs dataset. *CoRR*, abs/1907.01041, 2019. URL <http://arxiv.org/abs/1907.01041>.
- [49] Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel R. Bowman, Holger Schwenk, and Veselin Stoyanov. XNLI: evaluating cross-lingual sentence representations. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun’ichi Tsujii, editors, *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 2475–2485. Association for Computational Linguistics, 2018. doi: 10.18653/v1/d18-1269. URL <https://doi.org/10.18653/v1/d18-1269>.
- [50] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. In *5th Workshop on Energy Efficient Machine Learning and Cognitive Computing*. Neural Information Processing Systems, January 2019. URL https://openreview.net/forum?id=1u1I_xmPJLx.
- [51] Matt Gardner, Joel Grus, Mark Neumann, Oyvind Tafjord, Pradeep Dasigi, Nelson F. Liu, Matthew Peters, Michael Schmitz, and Luke Zettlemoyer. AllenNLP: A deep semantic natural language processing platform. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 1–6, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/W18-2501. URL <https://aclanthology.org/W18-2501>.
- [52] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, and Jamie Brew. Huggingface’s transformers: State-of-the-art natural language processing. *CoRR*, abs/1910.03771, 2019. URL <http://arxiv.org/abs/1910.03771>.

- [53] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- [54] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146, 2017. doi: 10.1162/tacl.a.00051. URL <https://aclanthology.org/Q17-1010>.
- [55] Phuon Thi Tran and Le Trieu Phong. On the convergence proof of amsgrad and a new version. *IEEE Access*, 7:61706–61716, 2019. ISSN 2169-3536. doi: 10.1109/access.2019.2916341. URL <http://dx.doi.org/10.1109/ACCESS.2019.2916341>.
- [56] Kingma, D.P., Ba, L.J., and Amsterdam Machine Learning lab (IVI, FNWI). Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations (ICLR)*. arXiv.org, 2015. URL [https://dare.uva.nl/personal/pure/en/publications/adam-a-method-for-stochastic-optimization\(a20791d3-1aff-464a-8544-268383c33a75\).html](https://dare.uva.nl/personal/pure/en/publications/adam-a-method-for-stochastic-optimization(a20791d3-1aff-464a-8544-268383c33a75).html).
- [57] Marcos Treviso and André F. T. Martins. The explanation game: Towards prediction explainability through sparse communication. In *Proceedings of the Third BlackboxNLP Workshop on Analyzing and Interpreting Neural Networks for NLP*, pages 107–118, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.blackboxnlp-1.10. URL <https://aclanthology.org/2020.blackboxnlp-1.10>.
- [58] Xiaobing Sun and Wei Lu. Understanding attention for text classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3418–3428, Online, July 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.acl-main.312. URL <https://aclanthology.org/2020.acl-main.312>.