# Is Attention Interpretable? 

Sofia Serrano* Noah A. Smith* ${ }^{\dagger}$<br>*Paul G. Allen School of Computer Science \& Engineering, University of Washington, Seattle, WA, USA<br>${ }^{\dagger}$ Allen Institute for Artificial Intelligence, Seattle, WA, USA<br>\{sofias6,nasmith\}@cs.washington.edu


#### Abstract

Attention mechanisms have recently boosted performance on a range of NLP tasks. Because attention layers explicitly weight input components' representations, it is also often assumed that attention can be used to identify information that models found important (e.g., specific contextualized word tokens). We test whether that assumption holds by manipulating attention weights in already-trained text classification models and analyzing the resulting differences in their predictions. While we observe some ways in which higher attention weights correlate with greater impact on model pre- dictions, we also find many ways in which this does not hold, i.e., where gradient-based rankings of attention weights better predict their effects than their magnitudes. We conclude that while attention noisily predicts input components' overall importance to a model, it is by no means a fail-safe indicator. ${ }^{1}$


## 1 Introduction

Interpretability is a pressing concern for many current NLP models. As they become increasingly complex and learn decision-making functions from data, ensuring our ability to understand why a particular decision occurred is critical.

Part of that development has been the incorporation of attention mechanisms (Bahdanau et al., 2015) into models for a variety of tasks. For many different problems-to name a few, machine translation (Luong et al., 2015), syntactic parsing (Vinyals et al., 2015), reading comprehension (Hermann et al., 2015), and language modeling (Liu and Lapata, 2018)-incorporating attention mechanisms into models has proven beneficial for performance. While there are many variants

[^0]of attention (Vaswani et al., 2017), each formulation consists of the same highlevel goal: calculating nonnegative weights for each input component (e.g., word) that together sum to 1 , multiplying those weights by their corresponding representations, and summing the resulting vectors into a single fixed-length representation.

Since attention calculates a distribution over inputs, prior work has used attention as a tool for interpretation of model decisions (Wang et al., 2016; Lee et al., 2017; Lin et al., 2017; Ghaeini et al., 2018). The existence of so much work on visualizing attention weights is a testament to attention's popularity in this regard; to name just a few examples of these weights being examined to understand a model, recent work has focused on goals from explaining and debugging the current system's decision (Lee et al., 2017; Ding et al., 2017) to distilling important traits of a dataset (Yang et al., 2017; Habernal et al., 2018).

Despite this, existing work on interpretability is only beginning to assess what computed attention weights actually communicate. In an independent and contemporaneous study, Jain and Wallace (2019) explore whether attention mechanisms can identify the relative importance of inputs to the full model, finding them to be highly inconsistent predictors. In this work, we apply a different analysis based on intermediate representation erasure to assess whether attention weights can in- stead be relied upon to explain the relative importance of the inputs to the attention layer itself. We find similar cause for concern: attention weights are only noisy predictors of even intermediate components' importance, and should not be treated as justification for a decision.

## 2 Testing for Informative Interpretability

We focus on five- and ten-class text classification models incorporating attention, as explaining the reasons for text classification has been a particular area of interest for recent work in interpretability (Yang et al., 2016; Ribeiro et al., 2016; Lei et al., 2016; Feng et al., 2018).

In order for a model to be interpretable, it must not only suggest explanations that make sense to people, but also ensure that those explanations accurately represent the true reasons for the model's decision. Note that this type of analysis does not rely on the true labels of the data; if a model produces an
incorrect output, but a faithful explanation for which factors were important in that calculation, we still consider it interpretable.

We take the implied explanation provided by visualizing attention weights to be a ranking of importance of the attention layer's input representations, which we denote $\mathcal{J}$ : if the attention allocated to item $i \in \mathcal{J}$ is higher than that allocated to item $j \in \mathcal{J}$, then $i$ is presumed "more important" than $j$ to the model's output. In this work, we focus on whether the attention weights' suggested importance ranking of $\mathcal{J}$ faithfully describes why the model produced its output, echoing existing work on explanation brittleness for other model components (Ghorbani et al., 2017).

### 2.1 Intermediate Representation Erasure



## Importance calculated from change in output

Figure 1: Our method for calculating the importance of representations corresponding to zeroed-out attention weights, in a hypothetical setting with four output classes.

We are interested in the impact of some contextualized inputs to an attention layer, $\mathcal{J}^{\prime} \subset \mathcal{J}$, on the model's output. To examine the importance of $\mathcal{J}^{\prime}$, we run the classification layer of the model twice (Figure 1): once without any modification, and once after renormalizing the attention distribution with $\mathcal{J}^{\prime \prime}$ s attention weights zeroed out, similar to other erasure-based work (Li et al., 2016; Feng et al., 2018). We then observe the resulting effects on the model's output. We erase at the attention layer to isolate the effects of the attention layer from the encoder preceding it. Our reasoning behind renormalizing is to keep the output document representation from artificially shrinking closer to 0 in a way never encountered during training, which could make subsequent measurements unrepresentative of the model's behavior in spaces to which it does map inputs.

One point worth noting is the facet of interpretability that our tests are designed to capture. By examining only how well attention represents the importance of intermediate quantities, which may themselves already have changed uninterpretably from the model's inputs, we are testing for a relatively low level of interpretability. So far, other work looking at attention has examined whether attention suffices as a holistic explanation for a model's decision (Jain and Wallace, 2019), which is a higher bar. We instead focus on the lowest standard of interpretability that attention might be expected to meet, ignoring prior model layers.

We denote the output distributions (over labels) as $\boldsymbol{p}$ (the original) and $\boldsymbol{q}_{\mathcal{J}^{\prime}}$ (where we erase attention for $\mathcal{J}^{\prime}$ ). The question now becomes how to operationalize "importance" given $\boldsymbol{p}$ and $\boldsymbol{q}_{\boldsymbol{J}^{\prime}}$. There are many quantities that could arguably capture information about importance. We focus on two: the Jensen-Shannon (JS) divergence between output distributions $\boldsymbol{p}$ and $\boldsymbol{q}_{\boldsymbol{J}^{\prime}}$, and whether the argmaxes of $\boldsymbol{p}$ and $\boldsymbol{q}_{\boldsymbol{\gamma}^{\prime}}$ differ, indicating a decision flip.

## 3 Data and Models

We investigate four model architectures on a topic classification dataset (Yahoo Answers; Zhang et al., 2015) and on three review ratings datasets: IMDB (Diao et
al., 2014), ${ }^{2}$ Yelp 2017, ${ }^{3}$ and Amazon (Zhang et al., 2015). Statistics for each dataset are listed in Table 1.

| Dataset | Average <br> Number of <br> Words (s.d.) | Average <br> Number of <br> Sentences <br> (s.d.) | Number of <br> Train + Dev <br> Instances | Number of <br> Test <br> Instances | Number of <br> Classes |
| :---: | :---: | :---: | :---: | :---: | :---: |
| Yahoo <br> Answers | $104(114)$ | $6.2(5.9)$ | $1,400,000$ | 50,000 | 10 |
| IMDB | $395(259)$ | $16.2(10.7)$ | 122,121 | 13,548 | 10 |
| Amazon | $73(48)$ | $4.3(2.6)$ | $3,000,000$ | 650,000 | 5 |
| Yelp | $125(109)$ | $7.0(5.6)$ | 650,000 | 50,000 | 5 |

Table 1: Dataset statistics.


Figure 2: Flat attention network (FLAN) demonstrating a convolutional encoder. Each contextualized word representation is the concatenation of two sizes of convolutions: one applied over the input representation and its two neighbors to either side, and the other applied over the input representation and its single neighbor to either side. For details, see Appendix A.1.

[^1]Our model architectures are inspired by the hierarchical attention network (HAN; Yang et al., 2016), a text classification model with two layers of attention, first to the word tokens in each sentence and then to the resulting sentence representations. The layer that classifies the document representation is linear with a softmax at the end.

We conduct our tests on the softmax formulation of attention, ${ }^{4}$ which is used by most models, including the HAN. Specifically, we use the additive formulation originally defined in Bahdanau et al. (2015). Given attention layer $\ell^{\prime}$ 's learned parameters, element $i$ of a sequence, and its encoded representation $\mathbf{h}_{i}$, the attention weight $\alpha_{i}$ is computed using $\ell$ 's learned context vector $\mathbf{c}_{\ell}$ as follows:

$$
\begin{gathered}
\mathbf{u}_{i}=\tanh \left(\mathbf{W}_{\ell} \mathbf{h}_{i}+\mathbf{b}_{\ell}\right) \\
\alpha_{i}=\frac{\exp \left(\mathbf{u}_{i}^{\top} \mathbf{c}_{\ell}\right)}{\sum_{j} \exp \left(\mathbf{u}_{j}^{\top} \mathbf{c}_{\ell}\right)}
\end{gathered}
$$

We evaluate on the original HAN architecture, but we also vary that architecture in two key ways:

1. Number of attention layers: besides exploring models with a final layer of attention over sentence representations (which we denote with a "HAN" prefix), we also train "flat" attention networks with only one layer of attention over all contextualized word tokens (which we denote with a "FLAN" prefix). In either case, though, we only run tests on models' final layer of attention.
2. Reach of encoder contextualization: The original HAN uses recurrent encoders to contextualize input tokens prior to an attention layer (specifically, bidirectional GRUs running over the full sequence). Aside from biRNNs, we also experiment with models that instead contextualize word vectors by convolutions on only a token's close neighbors, inspired by Kim

[^2](2014). See Figure 2 for a diagram of the FLAN architecture using a convolutional encoder. We denote this variant of an architecture with a "conv" suffix. Finally, we also test models that are trained with no contextualizing encoder at all; we denote these with a "noenc" suffix.

The classification accuracy of each of our trained models is listed in Table 3 in the appendix, along with training details for the different models.

## 4 Single Attention Weights' Importance

As a starting point for our tests, we investigate the relative importance of attention weights when only one weight is removed. Let $i^{*} \in \mathcal{J}$ be the component with the highest attention and let $\alpha_{i^{*}}$ be its attention. We compare $i^{* \prime}$ s importance to some other attended item's importance in two ways.

### 4.1 JS Divergence of Model Output Distributions

We wish to compare how $i^{* \prime}$ impact on the model's output distribution compares to the impact corresponding to a random attended item $r$ drawn uniformly from $\mathcal{J}$. Our first approach to this will be to calculate two JS divergences-one being the JS divergence of the model's original output distribution from its output distribution after removing only $i^{*}$, and the other after removing only $r$-and compare them to each other. We subtract the output JS divergence after removing $r$ from the output JS divergence after removing $i^{*}$ :

$$
\Delta \mathrm{JS}=\mathrm{JS}\left(\boldsymbol{p}, \boldsymbol{q}_{\left\{i^{*}\right\}}\right)-\mathrm{JS}\left(\boldsymbol{p}, \boldsymbol{q}_{\{r\}}\right)
$$

Equation 1
We plot this quantity against the difference $\Delta \alpha=\alpha_{i^{*}}-\alpha_{r}$ in Figure 3. We show results on the HANrnn, as the trends for the other models are very similar; see Figures $7-8$ and the tables in Figure 9 in the Appendix for full results.


Figure 3: Difference in attention weight magnitudes versus $\Delta J$ S for HANrnns, comparable to results for the other architectures; for their plots, see Appendix A. 2 .

## Counts of negative $\Delta \mathrm{JS}$ by difference in attention weight for HANrnns



Figure 4: These are the counts of test instances for the HANrnn models for which $i^{*}$ 's JS divergence was smaller, binned by $\Delta \alpha$. These counts comprise a small fraction of the test set sizes listed in Table 1.

Intuitively, if $i^{*}$ is truly the most important, then we would expect Eq. 1 to be positive, and that is what we find the vast majority of the time. In addition, examining Figure 3, we see that nearly all negative $\Delta J S$ values are close to 0 . By binning occurrences of negative $\Delta \mathrm{JS}$ values by the difference between $\alpha_{i^{*}}$ and $\alpha_{r}$ in Figure 4, we also see that in the cases where $i^{*}$ had a smaller effect, the gap between $i^{* \prime \prime}$ s attention and $r$ 's tends to be small. This is encouraging, indicating that in these cases, $i^{*}$ and $r$ are nearly "tied" in attention.

However, the picture of attention's interpretability grows somewhat more murky when we begin to consider the magnitudes of positive $\Delta J$ S values in Figure 3. We notice across datasets that even for quite large differences in attention weights like 0.4 , many of the positive $\Delta J$ S are still quite close to zero. Although we do
finally see an upward swing in $\Delta \mathrm{J} S$ values once $\Delta \alpha$ gets even larger, indicating only one very high attention weight in the distribution, this still leaves many open questions about exactly how much difference in impact $i^{*}$ and $r$ can typically be expected to have.

### 4.2 Decision Flips Caused by Zeroing Attention

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.5 | 8.7 |
|  | No | 1.3 | 89.6 |

Table 2a: Yahoo Answers, HANrnn

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 2.7 | 7.6 |
|  | No | 2.7 | 87.1 |

Table 2c: Amazon, HANrnn

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 2.2 | 12.2 |
|  | No | 1.4 | 84.2 |

Table 2b: IMDB, HANrnn

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 1.5 | 8.9 |
|  | No | 1.9 | 87.7 |

Table 2d: Yelp, HANrnn
Table 2: Percent of test instances in each decision-flip indicator variable category for each HANrnn.

Since attention weights are often interpreted as an explanation for a model's argmax decision, our second test looks at another, more immediately visible change in model outputs: decision flips. For clarity, we limit our discussion to results for the HANrnns, which reflect the same patterns observed for the other architectures. (Results for all other models are in Appendix A.2.)

Table 2 shows, for each dataset, a contingency table for the two binary random variables (i) does zeroing $\alpha_{i^{*}}$ (and renormalizing) result in a decision flip? and (ii) does doing the same for a different randomly chosen weight $\alpha_{r}$ result in a decision flip? To assess the comparative importance of $i^{*}$ and $r$, we consider when exactly one erasure changes the decision (off-diagonal cells). For attention to be interpretable, the blue, upper-right values ( $i^{*}$, not $r$, flips a decision) should be much larger than the orange, lower-left values ( $r$, not $i^{*}$, flips a decision), which should be close to zero. ${ }^{5}$

Although for some datasets in Table 2, the "orange" values (the lower-left values, where $r$, not $i^{*}$, flips a decision), are non-negligible, we mostly see that their fraction of total off-diagonal values mirrors the fraction of negative occurrences of Eq. 1 in Figure 4. However, it's somewhat startling that in the vast majority of cases, erasing $i^{*}$ does not change the decision ("no" row of each table). This is likely explained in part by the signal pertinent to the classification being distributed across a document (e.g., a "Sports" question in the Yahoo Answers dataset could signal "sports" in a few sentences, any one of which suffices to correctly categorize it). However, given that these results are for the HAN models, which typically compute attention over ten or fewer sentences, this is surprising.

Altogether, examining importance from a single-weight angle paints a tentatively positive picture of attention's interpretability, but also raises several questions about the many cases where the difference in impacts between $i^{*}$ and $r$ is almost identical (i.e., $\Delta \mathrm{J}$ S values close to 0 or the many cases where neither $i^{*}$ nor $r$ cause a decision flip). To answer these questions, we require tests with a broader scope.

## 5 Importance of Sets of Attention Weights

Often, we care about determining the collective importance of a set of components $\mathcal{J}^{\prime}$. To address that aspect of attention's interpretability and close gaps left by single-weight tests, we introduce tests to determine how multiple attention weights perform together as importance predictors.

[^3]
### 5.1 Multi-Weight Tests

For a hypothesized ranking of importance, such as that implied by attention weights, we would expect the items at the top of that ranking to function as a concise ex- planation for the model's decision. The less concise these explanations get, and the farther down the ranking that the items truly driving the model's decision fall, the less likely it becomes for that ranking to truly describe importance. In other words, we expect that the top items in a truly useful ranking of importance would comprise a minimal necessary set of information for making the model's decision.

The idea of a minimal set of inputs necessary to uphold a decision is not new; Li et al. (2016) use reinforcement learning to attempt to construct such a minimal set of words, Lei et al. (2016) train an encoder to constrain the input prior to classification, and much of the work that has been done on extractive summarization takes this concept as a starting point (Lin and Bilmes, 2011). However, such work has focused on approximating minimal sets, instead of evaluating the ability of other importance-determining "shortcuts" (such as attention weight orderings) to identify them. Nguyen (2018) leveraged the idea of minimal sets in a much more similar way to our work, comparing different input importance orderings.

Concretely, to assess the validity of an importance ranking method (e.g., attention), we begin erasing representations from the top of the ranking downward until the model's decision changes. Ideally, we would then enumerate all possible sub- sets of that instance's components, observe whether the model's decision changed in response to removing each subset, and then report whether the size of the minimal decision-flipping subset was equal to the number of items that had needed to be removed to achieve a decision flip by following the ranking. However, the exponential number of subsets for any given instance's sequence of components (word or sentence representations, in our case) makes such a strategy computationally prohibitive, and so we adopt a different approach.

Instead, in addition to our hypothesized importance ranking (attention weights), we consider alternative rankings of importance; if, using those, we repeatedly discover cases where removing a smaller subset of items would have sufficed to
change the decision, this signals that our candidate ranking is a poor indicator of importance.

### 5.2 Alternative Importance Rankings

Exhaustively searching the space of component subsets would be far too timeconsuming in practice, so we introduce three other ranking schemes.

The first is to randomly rank importance. We expect that this ranking will perform quite poorly, but it provides a point of comparison by which to validate that ranking by descending attention weights is at least somewhat informative.

The second ranking scheme, inspired by Li et al. (2015) and Feng et al. (2018), is to order the attention weights by the gradient of the decision function with respect to each calculated attention weight, in descending order. Since each of the datasets on which we evaluate has either five or ten output classes, we take the decision function given a real-valued model output vector to be

$$
d(\boldsymbol{x})=\frac{\exp \left(\max _{i}\left(\boldsymbol{x}_{i}\right)\right)}{\sum_{i} \exp \boldsymbol{x}_{i}}
$$

Unlike the last two proposed rankings, our third ranking scheme uses attention weights, but supplements them with information about the gradient. For this ranking, we multiply each of our calculated gradients from our previous proposed ranking scheme by their corresponding attention weight magnitude. Under this ordering, attended items that have both a high attention weight and a high calculated gradient with respect to their attention weight will be ranked most important.

We introduce these last two rankings as an attempt to discover smaller sets not produced by the attention weight ranking. Note, however, that we still do not take either as a gold-standard indicator of importance to the model, as with the gradient in Ross et al. (2017) and Melis and Jaakkola (2018), but merely as an alternative ordering method. The "gold standard" in our case would be the minimal set of attention weights to zero out for the decision to change, which none of our ordering methods will necessarily find for a particular instance.

### 5.3 Instances Excluded from Analysis

In cases where removing all but one input to the attention layer still does not produce a decision flip, we finish the process of removing components by removing the final representation and replacing the output of the attention layer with an arbitrary vector; we use the zero vector for our tests. Even so, since every real-valued vector output by the attention layer is mapped to an output distribution, removing this final item will still not change the classification decision for instances that the model happened to originally map to that same class. We exclude such instances for which the decision never changed from all subsequent analyses.

Fractions of Original Attention Weights Removed Before First Decision Flip Under Different Importance Rankings

HANrnns


Figure 5: The distribution of fractions of items removed before first decision flips on three model architectures under different ranking schemes. Boxplot whiskers represent the highest/lowest data point within 1.5 IQR of the higher/lower quartile, and dataset names at the bottom apply to their whole column. In several of the plots, the median or lower quartile aren't visible; in these cases, the median/lower quartile is either 1 or very close to 1 .

We also set aside any test instances with a sequence length of 1 for their final attention layer, as there is only one possible ordering for such cases.

### 5.4 Attention Does Not Optimally Describe Model Decisions

Examining our results in Figure 5, we immediately see that ranking importance by descending attention weights is not optimal for our models with encoders. While removing intermediate representations in decreasing order by attention weights often leads to a decision flip faster than a random ranking, it also clearly falls short of matching (or even approaching) the decision-flipping efficiency of either the gradient ordering or gradient-attention-product ordering in many cases.

In addition, although the product-based ranking often (but not always) requires slightly fewer removed items than the gradient ranking, we see that the purely gradient-based ranking ignoring attention magnitudes comes quite close to it, far outperforming purely attention-based orderings. For ten of our 16 models with encoders, removing by gradient found a smaller decision-flipping set of items than attention for over $50 \%$ of instances in that model's test set, with that percentage often being much higher. In fact, for every model with an encoder that we tested, there were at least 1.6 times as many test instances where the purely gradient-based ranking managed a decision flip faster than the attentionbased ranking than vice versa.

We do not claim that ranking importance by either descending gradients or descending gradient-attention products is optimal, but in many cases they discover much smaller decision-flipping sets of items than attention weights. Therefore, ranking representations in descending order by attention weight clearly fails to uncover a minimal set of decision-flipping information much of the time, which is a warning sign that we should be skeptical of trusting groups of attention weight magnitudes as importance indicators.

### 5.5 Decision Flips Often Occur Late

For all ordering schemes we tried, we were struck by the large fraction of items that had to be removed to achieve a decision flip in many models. This is slightly less surprising for the HANs, as they compute attention over shorter sequences of sentences (see Table 1). For the FLAN models, though, this result is highly unexpected. The sequences across which FLANs compute attention are usually
hundreds of tokens in length, meaning most attention weights will likely be minuscule.

The distributions of tokens removed by our different orderings that we see for the FLANrnns in Figure 5 are therefore remarkably high, especially given that all of our classification tasks have at least five output classes. We also note that due to the exponential nature of the softmax, softmax attention distributions typically contain only a few high-weighted items before the calculated weights become quite small, which can be misleading. In many cases, flipping the model's original decision requires digging deep into the small attention weights, with the highweighted components not actually being the reason for the decision.

For several of our models, especially the FLANs (which typically compute attention over hundreds of tokens), this fact is concerning from an explainability perspective. Lipton (2016) describes a model as "transparent" if "a person can contemplate the entire model at once." Applying this insight to the explanations suggested by attention, if an explanation rests on simultaneously considering hundreds of attended tokens necessary for a decision- even if that set were minimal-that would still raise serious transparency concerns.

### 5.6 Effects of Contextualization Scope on Attention's Interpretability

One last question we consider is whether the large number of items that are removed before decision flips can be explained in part by the scope of each model's contextualization. In machine translation, prior work has observed that recurrent encoders over a full sequence can "shift" tokens' signal in ways that cause subsequent attention layers to compute unintuitive off-by-one alignments (Koehn and Knowles, 2017). We hypothesize that in our text classification setting, the bidirectional recurrent structure of the HANrnn and FLANrnn encoders might instead be redistributing operative signal from a few informative input tokens across many others' contextualized representations. Under Different Importance Rankings


Figure 6: The distribution of fractions of items removed before decision flips on the encoderless model architectures under different ranking schemes. The Amazon FLANnoenc results have a very long tail; using the legend's order of rankings, the percentage of test instances with a removed fraction above 0.50 for that model is $12.4 \%, 2.8 \%, 0.9 \%$, and $0.5 \%$, respectively.

Comparing the decision flip results for the FLANconvs in Figure 5 to those for the FLANrnns supports this theory. We notice decision flips happening much faster than for either of the RNN-based model architectures, indicating that the biRNN effectively does learn to widely redistribute the classification signal. In contrast, the convolutional encoders only allow contextualization with respect to an input token's two neighbors to either side. We see similar results when comparing the two HAN architectures, albeit much more weakly (see Figure 10 in Appendix A.2); this is likely due to the smaller number of tokens being contextualized by the HANs (sentence representations instead of words), so that contextualization with respect to a token's close neighbors encompasses a much larger fraction of the full sequence.

We see this difference even more strongly when we compare to the encoderless model architectures, as shown in Figure 6. Compared to both other model architectures, we see the fraction of necessary items to erase for flipping the decision plummet. We also see random orderings mostly do better than before,
indicating more brittle decision boundaries, especially on the Amazon dataset. ${ }^{6}$ In this situation, we see attention magnitudes generally indicate importance on par with (or better than) gradients, but that the product-based ordering is still often a more efficient explanation.

While these differences themselves are not an argument against attention's interpretability, they highlight the distinction between attention's weighting of inter- mediate, contextualized representations and the model's use of the original input tokens themselves. Our RNN-based models' ability to maintain their original decision well past the point at which models using only local or no context have lost the signal driving their original decisions confirms that attention weights for a contextualized representation do not necessarily map neatly back to the original tokens. This might at least partly explain the striking near-indifference of the model's decision to the contributions of particular contextualized representations in both our RNN-based models and in Jain and Wallace (2019), who also use recurrent encoders.

However, the results from almost all models continue to support that ranking importance by attention is still not optimal; our non-random alternative rankings still uncover many cases where fewer items could be removed to achieve a decision flip than the attention weights imply.

## 6 Limitations

There are important limitations to the work we describe here, perhaps the most important of which is our focus on text classification. By choosing to focus on this task, we use the fact that decision flips are often not trivially achieved to ground our judgments of importance in model decision changes. However, for a task with a much larger output space (such as language modeling or machine translation) where almost anything might flip the decision, decision flips are likely too coarse a

[^4]signal to identify meaningful differences. Determining an analogous informative threshold in changes to model outputs would be key to expanding this sort of analysis to other groups of models.

A related limitation is our reliance in many of these tests on a fairly strict definition of importance tied to the output's argmax; an alternative definition of importance might assert that the highest attention weights should identify the most influential representations in pushing towards any output class, not just the argmax. Two of the core challenges that would need to be solved to test for how well attention meets this relaxed criterion would be meaningfully evaluating a single at- tended item's "importance" to multiple output classes for comparison to other at- tended items and, once again, determining what would truly indicate being "most influential" in the absence of decision flips as a guide to the output space.

Also, while we explore several model architectures in this work, there exist other attention functions such as multi-headed and scaled dot-product (Vaswani et al., 2017), as well as cases where a single attention layer is responsible for producing more than one attended representation, such as in self-attention (Cheng et al., 2016). These variants could have different interpretability properties. Like- wise, we only evaluate on final layers of attention here; in large models, lower-level layers of attention might learn to work differently.

## 7 Related and Future Work

We have adopted an erasure-based approach to probing the interpretability of computed attention weights, but there are many other possible approaches. For example, recent work has focused on which training instances (Koh and Liang, 2017) or which human-interpretable features were most relevant for a particular decision (Ribeiro et al., 2016; Arras et al., 2016). Others have explored alternative ways of comparing the behavior of proposed explanation methods (Adebayo et al., 2018). Yet another line of work focuses on aligning models with human feedback for what is interpretable (Fyshe et al., 2015; Subramanian et al., 2017), which could refine our idea of what defines a high-quality explanation derived from attention.

Finally, another direction for future work would be to extend the importanceranking comparisons that we deploy here for evaluation purposes into a method
for deriving better, more informative rankings, which in turn could be useful for the development of new, more interpretable models.

## 8 Conclusion

It is frequently assumed that attention is a tool for interpreting a model, but we find that attention does not necessarily correspond to importance. In some ways, the two correlate: comparing the highest attention weight to a lower weight, the high attention weight's impact on the model is often larger. However, the picture becomes bleaker when we consider the many cases where the highest attention weights fail to have a large impact. Examining these cases through multi-weight tests, we see that attention weights often fail to identify the sets of representations most important to the model's final decision. Even in cases when an attention- based importance ranking flips the model's decision faster than an alternative ranking, the number of zeroed attended items is often too large to be helpful as an explanation. We also see a marked effect of the contextualization scope preceding the attention layer on the number of attended items affecting the model's decision; while attention magnitudes do seem more helpful in uncontextualized cases, their lagging performance in retrieving decision rationales elsewhere is cause for concern. What is clear is that in the settings we have examined, attention is not an optimal method of identifying which attended elements are responsible for an output. Attention may yet be interpretable in other ways, but as an importance ranking, it fails to explain model decisions.

## Acknowledgments

This research was supported in part by a grant from the Allstate Corporation; findings do not necessarily represent the views of the sponsor. We thank R. Andrew Kreek, Paul Koester, Kourtney Traina, and Rebecca Jones for early conversations leading to this work. We also thank Omer Levy, Jesse Dodge, Sarthak Jain, Byron Wallace, and Dan Weld for helpful conversations, and our anonymous reviewers for their feedback.

## References

Julius Adebayo, Justin Gilmer, Michael Muelly, Ian Goodfellow, Moritz Hardt, and Been Kim. 2018. Sanity checks for saliency maps. In Advances in Neural Information Processing Systems.

Leila Arras, Franziska Horn, Grégoire Montavon, Klaus-Robert Müller, and Wojciech Samek. 2016. Explaining predictions of non-linear classifiers in NLP. arXiv preprint arXiv:1606.07298.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. Neural machine translation by jointly learning to align and translate. In Proceedings of the International Conference on Learning Representations.

Jianpeng Cheng, Li Dong, and Mirella Lapata. 2016. Long short-term memorynetworks for machine reading. In Proceedings of the Conference on Empirical Methods in Natural Language Processing.

Qiming Diao, Minghui Qiu, Chao-Yuan Wu, Alexander J. Smola, Jing Jiang, and Chong Wang. 2014. Jointly modeling aspects, ratings and sentiments for movie recommendation (JMARS). In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data mining.

Yanzhuo Ding, Yang Liu, Huanbo Luan, and Maosong Sun. 2017. Visualizing and understanding neural machine translation. In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers).

Shi Feng, Eric Wallace, Alvin Grissom II, Mohit lyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. Pathologies of neural models make interpretation difficult. In Proceedings of the Conference on Empirical Methods in Natural Language Processing.

Alona Fyshe, Leila Wehbe, Partha P. Talukdar, Brian Murphy, and Tom M. Mitchell. 2015. A compositional and interpretable semantic space. In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.

Reza Ghaeini, Xiaoli Z. Fern, and Prasad Tadepalli. 2018. Interpreting recurrent and attention-based neural models: a case study on natural language inference. arXiv preprint arXiv:1808.03894.

Amirata Ghorbani, Abubakar Abid, and James Zou. 2017. Interpretation of Neural Networks is Fragile. arXiv preprint arXiv:1710.10547.

Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. Before Name-calling: Dynamics and Triggers of Ad Hominem Fallacies in Web Argumentation. arXiv preprint arXiv:1802.06613.

Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In Advances in Neural Information Processing Systems.

Sarthak Jain and Byron C. Wallace. 2019. 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).

Yangfeng Ji and Noah A. Smith. 2017. Neural discourse structure for text categorization. In Proceedings of the Annual Meeting of the Association for Computational Linguistics.

Yoon Kim. 2014. Convolutional neural networks for sentence classification. In Proceedings of the Conference on Empirical Methods in Natural Language Processing.

Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Philipp Koehn and Rebecca Knowles. 2017. Six challenges for neural machine translation. In Proceedings of the First Workshop on Neural Machine Translation.

Pang Wei Koh and Percy Liang. 2017. Understanding black-box predictions via influence functions. arXiv preprint arXiv:1703.04730.

Jaesong Lee, Joong-Hwi Shin, and Jun-Seok Kim. 2017. Interactive Visualization and Manipulation of Attention-based Neural Machine Translation. In Proceedings
of the 2017 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 121-126.

Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. Rationalizing neural predictions. In Proceedings of the Conference on Empirical Methods in Natural Language Processing.

Jiwei Li, Xinlei Chen, Eduard Hovy, and Dan Jurafsky. 2015. Visualizing and understanding neural models in NLP. arXiv preprint arXiv:1506.01066.

Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. Understanding neural networks through representation erasure. arXiv preprint arXiv:1612.08220.

Hui Lin and Jeff Bilmes. 2011. A class of submodular functions for document summarization. In Proceedings of the Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1.

Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. arXiv preprint arXiv:1703.03130.

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

Yang Liu and Mirella Lapata. 2018. Learning structured text representations. Transactions of the Association of Computational Linguistics, 6:63-75.

Minh-Thang Luong, Hieu Pham, and Christopher D. Manning. 2015. Effective approaches to attention-based neural machine translation. arXiv preprint arXiv:1508.04025.

André Martins and Ramón Astudillo. 2016. From softmax to sparsemax: A sparse model of attention and multi-label classification. In International Conference on Machine Learning.

David Alvarez Melis and Tommi Jaakkola. 2018. Towards robust interpretability with self-explaining neural networks. In Advances in Neural Information Processing Systems.

Dong Nguyen. 2018. Comparing automatic and human evaluation of local explanations for text classification. In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers).

Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should I trust you?": Explaining the predictions of any classifier. In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data mining.

Andrew Slavin Ross, Michael C. Hughes, and Finale Doshi-Velez. 2017. Right for the right reasons: Training differentiable models by constraining their explanations. In Proceedings of the International Joint Conference on Artificial Intelligence.

Anant Subramanian, Danish Pruthi, Harsh Jhamtani, Taylor Berg-Kirkpatrick, and Eduard Hovy. 2017. SPINE: Sparse interpretable neural embeddings. arXiv preprint arXiv:1711.08792.

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 Advances in Neural Information Processing Systems.

Oriol Vinyals, Łukasz Kaiser, Terry Koo, Slav Petrov, Ilya Sutskever, and Geoffrey Hinton. 2015. Grammar as a foreign language. In Advances in Neural Information Processing Systems.

Yequan Wang, Minlie Huang, Li Zhao, et al. 2016. Attention-based LSTM for aspect-level sentiment classification. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, pages 606-615.

Fan Yang, Arjun Mukherjee, and Eduard Dragut. 2017. Satirical news detection and analysis using attention mechanism and linguistic features. arXiv preprint arXiv:1709.01189.

Zichao Yang, Diyi Yang, Chris Dyer, Xiaodong He, Alex Smola, and Eduard Hovy. 2016. Hierarchical attention networks for document classification. In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies.

Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. In Advances in Neural Information Processing Systems.

## A. 1 Model Hyperparameters and Performance

We lowercased all tokens during preprocessing and used all hyperparameters specified in (Yang et al., 2016), except for those related to the optimization algorithm or, in the case of the convolutional or no-encoder models, the encoder. For each convolutional encoder, we trained two convolutions: one sweeping over five to- kens, and one sweeping over three. As the output representation of token $x$, we then concatenated the outputs of the five-token and three-token convolutions centered on $x$. Unless otherwise noted, to train each model, we used Adam (Kingma and Ba, 2014) with gradient clipping of 10.0 and a patience value of 5 , so we would stop training a model if five epochs elapsed without any improvement in validation set accuracy. In addition, for each model, we specified a learning rate for training, and dropout before each encoder layer (or attention layer, for the encoderless models) and also within the classification layer. For the HAN models, these are the values we used:

- Yahoo Answers HANrnn, Yahoo Answers HANconv
- Pre-sentence-encoder dropout: 0.4445
- Pre-document-encoder dropout: 0.2202
- Classification layer dropout: 0.3749
- Learning rate: 0.0004
- IMDB HANrnn
- Pre-sentence-encoder dropout: 0.4445
- Pre-document-encoder dropout: 0.2202
- Classification layer dropout: 0.2457
- Learning rate: 0.0004
- Amazon HANrnn, Amazon HANconv
- Pre-sentence-encoder dropout: 0.6
- Pre-document-encoder dropout: 0.2
- Classification layer dropout: 0.4
- Learning rate: 0.0002
- Amazon HANnoenc
- Pre-sentence-encoder dropout: 0.6
- Pre-document-encoder dropout: 0.2
- Classification layer dropout: 0.4
- Learning rate: 0.0002
- Patience: 10
- Yelp HANrnn, Yelp HANconv
- Pre-sentence-encoder dropout: 0.7
- Pre-document-encoder dropout: 0.1
- Classification layer dropout: 0.7
- Learning rate: 0.0001
- Yelp HANnoenc
- Pre-sentence-encoder dropout: 0.7
- Pre-document-encoder dropout: 0.1
- Classification layer dropout: 0.7
- Learning rate: 0.0001
- Patience: 10
- Yahoo Answer HANnoenc
- Pre-sentence-encoder dropout: 0.4445
- Pre-document-encoder dropout: 0.2202
- Classification layer dropout: 0.3749
- Learning rate: 0.0004
- Patience: 10
- IMDB HANconv
- Pre-sentence-encoder dropout: 0.4445
- Pre-document-encoder dropout: 0.2202
- Classification layer dropout: 0.2457
- Learning rate: 0.0004
- IMDB HANnoenc
- Pre-sentence-encoder dropout: 0.4445
- Pre-document-encoder dropout: 0.2202
- Classification layer dropout: 0.2457
- Learning rate: 0.0004
- Patience: 10

For the FLAN models, these are the values we used:

- Yahoo Answers FLANrnn, Yahoo Answers FLANconv
- Pre-document-encoder dropout: 0.4445
- Classification layer dropout: 0.4457
- Learning rate: 0.0004
- IMDB FLANrnn, IMDB FLANconv
- Pre-document-encoder dropout: 0.4445
- Classification layer dropout: 0.3457
- Learning rate: 0.0004
- Amazon FLANrnn, Amazon FLANconv
- Pre-document-encoder dropout: 0.6
- Classification layer dropout: 0.4
- Learning rate: 0.0002
- Amazon FLANnoenc
- Pre-document-encoder dropout: 0.6
- Classification layer dropout: 0.4
- Learning rate: 0.0002
- Patience: 10
- Yelp FLANrnn, Yelp FLANconv
- Pre-document-encoder dropout: 0.7
- Classification layer dropout: 0.7
- Learning rate: 0.0001
- Yelp FLANnoenc
- Pre-document-encoder dropout: 0.7
- Classification layer dropout: 0.7
- Learning rate: 0.0001
- Patience: 10
- Yahoo Answers FLANnoenc
- Pre-document-encoder dropout: 0.4445
- Classification layer dropout: 0.4457
- Learning rate: 0.0004
- Patience: 10
- IMDB FLANnoenc
- Pre-document-encoder dropout: 0.4445
- Classification layer dropout: 0.3457
- Learning rate: 0.0004
- Patience: 10

| Dataset | HANrnn | HANconv | HANnoenc | FLANrnn | FLANconv | FLANnoenc |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Yahoo <br> Answers | 74.6 | 72.8 | 73.1 | 75.5 | 73.1 | 72.3 |
| IMDB | 50.3 | 48.9 | 46.1 | 49.1 | 48.2 | 45.4 |
| Amazon | 56.9 | 55.3 | 51.2 | 56.6 | 54.4 | 50.2 |
| Yelp | 63.0 | 61.0 | 58.6 | 62.3 | 60.7 | 58.2 |

Table 3: Classification accuracy of the different trained models on their respective test sets

Trained model classification accuracies are reported in Table 3. We note that our IMDB data and Yelp data are different sets of reviews from those used by Yang et al. (2016), so our reported performances are not directly comparable to theirs. We were unable to reach a comparable performance for the Amazon dataset (and Yelp dataset, although different) to that in (Yang et al., 2016). We suspect that this is due to not pretraining the word2vec embeddings used by the model for long enough, combined with memory limitations on our hardware that necessitated decreasing our batch size in many cases. However, as noted in section 3 , the analysis that we perform does not depend on model accuracy. It's also worth noting that for the datasets for which we are able to get results that either pass or come close to the accuracies listed in the original HAN paper, the patterns we see in the results for the tests that we run are the same as the patterns that we see for the others.

## A. 2 Full Sets of Plots

Here we include the full sets of result plots for all models for all tests we describe in the paper, in order of appearance.

In Figure 7, we see that the majority of $\Delta \mathrm{J} S$ values continue to fall above 0 , and that most are still close to 0 . One point not stated in the paper, though, is that the upswing in $\Delta \mathrm{J} S$ values as the difference between $i^{* \prime}$ s weight and a randomly chosen weight increases tends to occur slightly earlier for models with less contextualization, implying that the improving efficiency of the attention-based ranking at flipping the decision as contextualization scope shrinks is also reflected in single-weight test results.


Figure 7: Differences in attention weight magnitude plotted against $\Delta \mathrm{JS}$ for all datasets and architectures.


Figure 8: Counts of negative $\Delta \mathrm{JS}$ values grouped by the difference in their corresponding attention weights for all datasets and architectures.

Figure 9 (on the pages following this description): Using the definition of $i^{*}$ given in section 4 (the highest-attention-weight attended item) and comparing to a different randomly selected attended item, these were the percentages of test instances that fell into each decision-flip indicator variable category for each of the four test sets on all models. Since we require our random item not to be $i^{*}$, we exclude any instances with a final sequence length of 1 (one sentence for the HANs, one word for the FLANs) from analysis.

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.5 | 8.7 |
|  | No | 1.3 | 89.6 |

Figure 9, HANrnn table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 2.7 | 7.6 |
|  | No | 2.7 | 87.1 |

Figure 9, HANrnn table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.3 | 9.0 |
|  | No | 0.3 | 90.3 |

Figure 9, FLANrnn table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.6 | 7.3 |
|  | No | 0.5 | 91.6 |

Figure 9, FLANrnn table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 2.2 | 12.2 |
|  | No | 1.4 | 84.2 |

Figure 9, HANrnn table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 1.5 | 8.9 |
|  | No | 1.9 | 87.7 |

Figure 9, HANrnn table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ |  |  |  |
| Decision <br> flip? | Yes | 0.1 | 14.5 |
|  | No | 0.1 | 85.3 |

Figure 9, FLANrnn table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.3 | 8.0 |
|  | No | 0.3 | 91.4 |

Figure 9, FLANrnn table: Yelp

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 1.2 | 3.3 |
|  | No | 2.1 | 93.4 |

Figure 9, HANconv table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 4.4 | 13.1 |
|  | No | 5.4 | 77.0 |

Figure 9, HANconv table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.4 | 6.8 |
|  | No | 0.8 | 91.9 |

Figure 9, FLANconv table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 1.3 | 6.8 |
|  | No | 1.6 | 90.2 |

Figure 9, FLANconv table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 3.4 | 14.7 |
|  | No | 2.6 | 79.3 |

Figure 9, HANconv table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 3.3 | 11.1 |
|  | No | 4.0 | 81.6 |

Figure 9, HANconv table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.2 | 17.3 |
|  | No | 0.3 | 82.2 |

Figure 9, FLANconv table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ |  |  |  |
| Decision <br> flip? | Yes | 0.8 | 12.1 |
|  | No | 0.7 | 86.3 |

Figure 9, FLANconv table: Yelp

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 2.5 | 18.2 |
|  | No | 3.7 | 75.7 |

Figure 9, HANnoenc table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 13.8 | 25.8 |
|  | No | 6.0 | 54.3 |

Figure 9, HANnoenc table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.7 | 5.8 |
|  | No | 1.1 | 92.4 |

Figure 9, FLANnoenc table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 2.5 | 24.8 |
|  | No | 1.6 | 71.0 |

Figure 9, FLANnoenc table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 6.7 | 34.7 |
|  | No | 3.2 | 55.4 |

Figure 9, HANnoenc table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 8.4 | 18.0 |
|  | No | 5.2 | 68.4 |

Figure 9, HANnoenc table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 0.3 | 27.6 |
|  | No | 0.3 | 71.8 |

Figure 9, FLANnoenc table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i^{*}:$ <br> Decision <br> flip? | Yes | 1.4 | 14.1 |
|  | No | 1.3 | 83.2 |

Figure 9, FLANnoenc table: Yelp

Looking at where negative $\Delta J$ S values tend to occur in Figure 8, we once again see that they tend to cluster around cases where the difference between the highest and randomly chosen attention weights is close to 0 . There are some exceptions, however; perhaps the most obvious are the fat tails of these counts for the Yahoo Answer HAN models. Considering the highest-attention-weight ranking of importance for all Yahoo Answers HAN models in Figure 10 struggle in flipping the decision quickly, it may be that attention is less helpful than usual in identifying importance in its case, which could explain this discrepancy.

In Figure 9, we list contingency tables for all $i^{*}$-versus-random single-weight decision-flip tests. We continue to see higher values overall in our blue cells (upper-right values, where $i^{*}$, not $r$, flips a decision) than orange (lower-left values, where $r$, not $i^{*}$, flips a decision), as described in section 4.2. The most general change we notice across all the tables is that in the encoderless case, there are more test instances (often many more) where at least one of $i^{*}$ or our random attended item flipped the decision than for any other architecture, except in the case of the Yahoo Answers FLAN. Thinking about why this might be, we recall that in the encoderless case, word embeddings are much more directly responsible for encoding a decision. Yahoo Answers is our only topic classification dataset, where keywords like "computer" or "basketball" might be much clearer indicators of a topic than, say, "like" or "love" would be indicators of a rating of 8 versus 9 . This likely leads to much less certain decisions being encoded in the word embeddings of the non-Yahoo Answers datasets. For all other models, and in the case where potentially contradictory Yahoo Answers word embeddings are blended together before the final layer of attention (its HANnoenc), it is likely that decisions are simply more brittle overall.

Finally, in Figure 10, we include the full set of fraction-removed distributions for the first decision flips reached under the different rankings we explored.

## A. 3 Additional Tests

Besides the tests we describe in the main paper, some of the other tests that we ran provide additional insights into our results. We briefly describe those here.

In Figure 11, we provide the distributions of the original attention probability distributions that were zeroed at the point when different ranking schemes achieved their first decision flips. (Equivalently, these are the distributions of the sums of the zeroed attention weights described in Figure 10, only without repeated normalization.) We include these results to give a sense of which attention magnitudes the different rankings typically place towards the top. We notice that this probability mass required to change a decision is often quite high, which is unsurprising for the attention-based ranking, given that it frequently requires removing many items to flip decisions and attention distributions tend to have just a few high weights.

Besides that, the main takeaway that we see here is that for most models, the distribution of attention probability masses zeroed by our gradient-based ranking or our product-based ranking is often shifted down by around 0.25 or more compared to the corresponding attention probability mass distribution for the attention-based ranking, which is a fairly large difference. This would seem to imply that these alternative rankings (which usually flip decisions faster) tend to differ in relatively substantial ways from the rankings suggested by the pure attention weights, not just in the long tail of their orderings, which is another warning sign against attention's interpretability.

The final set of tests that we include in Figures 12 and 13 consist of rerunning our single-weight decision-flip tests on the single "most important" attention weights in their respective attention distributions as suggested by our alternative rankings (gradient-based and product-based rankings) instead of attention magnitudes. These results serve two functions: first, they imply still more information about when the top weight suggested by an alternative, faster-decision-flipping ranking differs from the top attention weight. Intuitively, if we observe large differences between the sum of the "yes" row for one contingency table and the "yes" rows for the other rankings' tables on that same model, this is likely due to differences in the frequencies with which the highest-ranked items achieve a decision flip, indicating differences in highest-ranked items ("likely" because of the noise added by the random sampling).

The second piece of information that these tests provide is a lower bound (via the sum of the "yes" rows) for the number of cases where rankings flip a decision as quickly as possible (i.e., in the first item). For context, the sum of the "yes" row is higher than the corresponding sum in Figure 9 for all contingency tables using our
product-based ordering. For the gradient-based ordering, however, this sum is actually lower than for the attention-based ranking's tables in 14 out of our 24 models. This tells us that our gradient-based method often finds fewer singleitem ways of flipping decisions than the attention-based ranking, so in order to achieve its more efficient overall distribution of flips that we see for many models in Figure 10, it must usually flip decisions faster than attention in cases where both its ranking and the attention-based ranking require multiple removed weights.

Fractions of Original Attention Weights Removed Before First Decision Flip Under Different Importance Rankings


Figure 10: Distribution of fraction of attention weights that had to be removed by different ranking schemes to change each model architecture's decisions for each of the four datasets. The different rankings (aside from "Attention", which corresponds to the attention weight magnitudes in descending order) are described in section 5.2.

Probability Masses of Original Attention Distributions Removed Before
First Decision Flip Under Different Importance Rankings


Figure 11: Distribution of probability masses that had to be removed by different ranking schemes to change each model architecture's decisions for each of the four datasets. While we do not discuss these in the paper due to space constraints, we notice that in most cases, a high fraction of the original attention distribution's probability mass must be zeroed before the (renormalized) modified attended representation results in a changed decision using the Attention ranking.

Figure 12 (on the pages following this description): Let $i_{g}^{*}$ be the highest-ranked attended item using our purely gradient-based ranking of importance described in section 5.2. We rerun our single-weight decision flip tests using this new $i_{g}^{*}$, comparing to a different randomly selected attended item. These were the percentages of test instances that fell into each decision-flip indicator variable category for each of the four test sets on all models. Since we require our random item not to be $i_{g}^{*}$, we exclude any instances with a final sequence length of 1 (one sentence for the HANs, one word for the FLANs) from analysis.

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 1.6 | 7.9 |
|  | No | 0.2 | 90.3 |

Figure 12, HANrnn table: Yahoo

| Amazon |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |  |
|  |  | Yes | No |  |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 4.9 | 10.2 |  |
|  | No | 0.4 | 84.5 |  |

Figure 12, HANrnn table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 0.5 | 7.2 |
|  | No | 0.2 | 92.2 |

Figure 12, FLANrnn table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 0.9 | 7.1 |
|  | No | 0.2 | 91.8 |

Figure 12, FLANrnn table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 2.8 | 13.7 |
|  | No | 0.6 | 82.9 |

Figure 12, HANrnn table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yos | 3.0 | 10.0 |
|  | No | 0.4 | 86.6 |

Figure 12, HANrnn table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}$ <br> Decision <br> flip? | Yes | 0.1 | 10.9 |
|  | No | 0.1 | 88.9 |

Figure 12, FLANrnn table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}$ : <br> Decision <br> flip? | Yes | 0.4 | 7.4 |
|  | No | 0.1 | 92.1 |

Figure 12, FLANrnn table: Yelp

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 2.8 | 3.0 |
|  | No | 0.3 | 93.8 |

Figure 12, HANconv table: Yahoo

| Amazon |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |  |
|  |  | Yes | No |  |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 9.1 | 16.7 |  |
|  | No | 0.8 | 73.5 |  |

Figure 12, HANconv table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 0.9 | 12.7 |
|  | No | 0.3 | 86.1 |

Figure 12, FLANconv table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 2.0 | 15.4 |
|  | No | 1.0 | 81.6 |

Figure 12, FLANconv table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | No | 1.0 |

Figure 12, HANconv table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yos | 6.2 | 13.2 |
|  | No | 1.0 | 79.6 |

Figure 12, HANconv table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 0.4 | 16.6 |
|  | No | 0.2 | 82.8 |

Figure 12, FLANconv table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 1.2 | 11.3 |
|  | No | 0.4 | 87.2 |

Figure 12, FLANconv table: Yelp

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 5.3 | 15.7 |
|  | No | 0.8 | 78.2 |

Figure 12, HANnoenc table: Yahoo

| Amazon |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |  |
|  |  | Yes | No |  |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 18.1 | 32.0 |  |
|  | No | 1.7 | 48.2 |  |

Figure 12, HANnoenc table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 0.7 | 4.1 |
|  | No | 1.1 | 94.1 |

Figure 12, FLANnoenc table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 2.9 | 17.0 |
|  | No | 1.2 | 78.8 |

Figure 12, FLANnoenc table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ | Yes | 7.4 | 30.5 |
| Decision <br> flip? | No | 2.9 | 59.3 |

Figure 12, HANnoenc table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yos | 11.8 | 26.2 |
|  | No | 1.6 | 60.4 |

Figure 12, HANnoenc table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yos | 0.3 | 11.1 |
|  | No | 88.2 |  |

Figure 12, FLANnoenc table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{g}^{*}:$ <br> Decision <br> flip? | Yes | 1.2 | 6.2 |
|  | No | 1.4 | 91.2 |

Figure 12, FLANnoenc table: Yelp

Figure 13 (on the pages following this description): Let $i_{p}^{*}$ be the highest-ranked attended item using our attention-gradient product ranking of importance described in section 5.2. Once again, we rerun our single-weight decision flip tests using this new $i_{p}^{*}$, comparing to a different randomly selected attended item. These were the percentages of test instances that fell into each decision-flip indicator variable category for each of the four test sets on all models. Since we require our random item not to be $i_{p}^{*}$, we exclude any instances with a final sequence length of 1 (one sentence for the HANs, one word for the FLANs) from analysis.

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 1.3 | 9.5 |
|  | No | 0.5 | 88.7 |

Figure 13, HANrnn table: Yahoo

| Amazon |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |  |
|  |  | Yes | No |  |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 4.5 | 10.5 |  |
|  | No | 0.8 | 84.2 |  |

Figure 13, HANrnn table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 0.6 | 13.2 |
|  | No | 0.0 | 86.2 |

Figure 13, FLANrnn table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 0.9 | 10.5 |
|  | No | 0.1 | 88.5 |

Figure 13, FLANrnn table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ | Yes | 2.9 | 15.5 |
| Decision <br> flip? | No | 0.6 | 81.0 |

Figure 13, HANrnn table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 2.7 | 11.3 |
|  | No | 0.8 | 85.3 |

Figure 13, HANrnn table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 0.2 | 18.3 |
|  | No | 0.1 | 81.5 |

Figure 13, FLANrnn table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 0.4 | 10.6 |
|  | No | 0.1 | 88.9 |

Figure 13, FLANrnn table: Yelp

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 2.8 | 5.1 |
|  | No | 0.4 | 91.7 |

Figure 13, HANconv table: Yahoo

| Amazon |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |  |
|  |  | Yes | No |  |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 8.8 | 18.2 |  |
|  | No | 1.0 | 72.0 |  |

Figure 13, HANconv table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 1.1 | 19.3 |
|  | No | 0.0 | 79.6 |

Figure 13, FLANconv table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 2.7 | 32.7 |
|  | No | 0.2 | 64.3 |

Figure 13, FLANconv table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 5.0 | 20.5 |
|  | No | 0.9 | 73.6 |

Figure 13, HANconv table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yos | 6.3 | 16.0 |
|  | No | 0.8 | 76.9 |

Figure 13, HANconv table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yo | 0.0 | 72.8 |
|  |  | Nes | 0.5 |

Figure 13, FLANconv table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 1.3 | 20.4 |
|  | No | 0.2 | 78.1 |

Figure 13, FLANconv table: Yelp

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 5.6 | 23.1 |
|  | No | 0.6 | 70.7 |

Figure 13, HANnoenc table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 18.4 | 33.9 |
|  | No | 1.5 | 46.2 |

Figure 13, HANnoenc table: Amazon

| Yahoo Answers |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 1.8 | 19.4 |
|  | No | 0.0 | 78.8 |

Figure 13, FLANnoenc table: Yahoo

| Amazon |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 3.8 | 35.9 |
|  | No | 0.3 | 60.0 |

Figure 13, FLANnoenc table: Amazon

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yes | 9.2 | 42.9 |
|  | No | 0.8 | 47.1 |

Figure 13, HANnoenc table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yos | 11.9 | 27.8 |
|  | No | 1.6 | 58.7 |

Figure 13, HANnoenc table: Yelp

| IMDB |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yos | 0.6 | 36.2 |
|  | No | 0.1 | 63.1 |

Figure 13, FLANnoenc table: IMDB

| Yelp |  |  |  |
| :---: | :---: | :---: | :---: |
|  |  | Remove random: <br> Decision flip? |  |
|  |  | Yes | No |
| Remove $i_{p}^{*}:$ <br> Decision <br> flip? | Yos | 2.3 | 26.3 |
|  | No | 0.3 | 71.1 |

Figure 13, FLANnoenc table: Yelp


[^0]:    ${ }^{1}$ Code is available at https://github.com/serrano-s/attn-tests.

[^1]:    ${ }^{2}$ downloaded from github.com/nihalb/JMARS
    ${ }^{3}$ from www.yelp.com/dataset challenge

[^2]:    ${ }^{4}$ Alternatives such as sparse attention (Martins and Astudillo, 2016) and unnormalized attention (Ji and Smith, 2017) have been proposed.

[^3]:    ${ }^{5}$ We see this pattern especially strongly for FLANs (see Appendix), which is unsurprising since $\mathcal{J}$ is all words in the input text, so most attention weights are very small.

[^4]:    ${ }^{6}$ This is likely due to the fact that with no contextualization, the final attended representations are just a linear combination of the input embeddings, so the embeddings themselves are responsible for learning to directly encode a decision. Since Amazon has the largest ratio of documents (which probably vary in their label) to unique word embeddings by a factor of more than two times any other dataset's, and the final attended representations in the FLANnoencs are unaggregated word embeddings, it stands to reason that the lack of encoders would be a much bigger obstacle in its case.

