

# Easy to Decide, Hard to Agree: Reducing Disagreements Between Saliency Methods

Josip Jukić<sup>♣,1</sup>, Martin Tutek<sup>♣,♡,2</sup>, Jan Šnajder<sup>1</sup>

<sup>1</sup>TakeLab, University of Zagreb   <sup>2</sup>UKP Lab, Technical University of Darmstadt  
{josip.jukic, jan.snajder}@fer.hr  
tutek@ukp.informatik.tu-darmstadt.de

## Abstract

A popular approach to unveiling the black box of neural NLP models is to leverage saliency methods, which assign scalar importance scores to each input component. A common practice for evaluating whether an interpretability method is *faithful* has been to use evaluation-by-agreement – if multiple methods agree on an explanation, its credibility increases. However, recent work has found that saliency methods exhibit weak rank correlations even when applied to the same model instance and advocated for alternative diagnostic methods. In our work, we demonstrate that rank correlation is sensitive to small perturbations when evaluating agreement and argue that Pearson- $r$  could be a better-suited alternative. We further show that regularization techniques that increase faithfulness of attention explanations also increase agreement between saliency methods. By connecting our findings to instance categories based on training dynamics, we show that the agreement of saliency method explanations is very low for easy-to-learn instances. Finally, we connect the improvement in agreement across instance categories to local representation space statistics of instances, paving the way for work on analyzing which intrinsic model properties improve their predisposition to interpretability methods.<sup>1</sup>

## 1 Introduction

Following the meteoric rise of the popularity of neural NLP models during the neural revolution, they have found practical usage across a plethora of domains and tasks. However, in a number of high-stakes domains such as law (Kehl and Kessler, 2017), finance (Grath et al., 2018), and medicine (Caruana et al., 2015), the opacity of deep learning

methods needs to be addressed. In the area of explainable artificial intelligence (XAI), one of the major recent efforts is to unveil the neural black box and produce explanations for the end-user. There are various approaches to rationalizing model predictions, such as using the attention mechanism (Bahdanau et al., 2014), saliency methods (Denil et al., 2014; Bach et al., 2015; Ribeiro et al., 2016; Lundberg and Lee, 2017; Shrikumar et al., 2017; Sundararajan et al., 2017), rationale generation by-design (Lei et al., 2016; Bastings et al., 2019; Jain et al., 2020), or self-rationalizing models (Marasovic et al., 2022). These methods have to simultaneously satisfy numerous desiderata to have practical application in high-stakes scenarios: they have to be *faithful* – an accurate representation of the inner reasoning process of the model, and *plausible* – convincing to human stakeholders.

When evaluating faithfulness in using attention as explanations, Jain and Wallace (2019) have shown that attention importance scores do not correlate well with gradient-based measures of feature importance. The authors state that although gradient-based measures of feature importance should not be taken as ground truth, one would still expect importance measures to be highly agreeable, bringing forth the *agreement-as-evaluation* paradigm (Abnar and Zuidema, 2020; Meister et al., 2021). While the imperfect agreement is something one could expect as interpretability methods differ in their formulation, and it is reasonable to observe differences in importance scores, subsequent work has shown that saliency methods exhibit low agreement scores even when applied to the *same model instance* (Neely et al., 2021). Since a single trained model instance can only have a single feature importance ranking for its decision, disagreement of saliency methods implies that at least one, if not all methods, do not produce faithful explanations – placing doubt on their practical relevance. It has been hypothesized that unfaithfulness of attention

<sup>♣</sup>Equal contribution

<sup>♡</sup>Work done while author was at TakeLab

<sup>1</sup>Our code is available at <https://github.com/mttk/unity-in-xai-court>

is caused by input entanglement in the hidden space (Jain and Wallace, 2019). This claim has later been experimentally verified through results showing that regularization techniques targeted to reduce entanglement significantly improve the faithfulness of attention-based explanations (Mohankumar et al., 2020; Tutek and Šnajder, 2020). While entanglement in the hidden space is clearly a problem in the case of attention explanations, where attention weights directly pertain to hidden states, we also hypothesize that representation entanglement could cause similar issues for gradient- and propagation-based explainability methods – which might not be able to adequately disentangle importance when propagating toward the inputs.

In our work, we first take a closer look at whether the rank correlation is an appropriate method for evaluating agreement and confirm that, as hypothesized in previous work, small differences in values of saliency scores significantly affect agreement scores. We argue that a linear correlation method such as Pearson- $r$  is less sensitive to perturbations since the exact ranking order of features is not as crucial for agreement as the relative importance values, which Pearson- $r$  adequately captures. We hypothesize that the cause of saliency method disagreements is rooted in representation entanglement and experimentally show that agreement can be significantly improved by regularization techniques such as tying (Tutek and Šnajder, 2020) and conicity (Mohankumar et al., 2020). The fact that regularization methods, which were originally aimed at improving faithfulness of attention, also improve agreement between saliency methods suggests that the two problems have the same underlying cause. Taking the analysis deeper, we apply techniques from dataset cartography (Swayamdipta et al., 2020) and show that, surprisingly, the explanations of easy-to-learn instances exhibit a lower agreement than of ambiguous instances. We further analyze how local curvature of the representation space morphs when regularization techniques are applied, paving the way for further analysis of (dis)agreements between interpretability methods.

## 2 Background and Related Work

Explainability methods come in different flavors determined by the method of computing feature importance scores. Saliency methods perform *post-hoc* analysis of the trained black-box model by either leveraging gradient information (Denil et al.,

2014; Sundararajan et al., 2017), modifying the backpropagation rules (Bach et al., 2015; Shrikumar et al., 2017), or training a shallow interpretable model to locally approximate behavior of the black-box model (Ribeiro et al., 2016), all with the goal of assigning scalar saliency scores to input features. Alternatively, if the analyzed model is capable of generating text, one can resort to self-rationalization by prompting the trained model to generate an explanation for its decision (Marasovic et al., 2022). In contrast to *post-hoc* explanations, *inherently interpretable* models produce explanations as part of their decision process, either by masking a proportion of input tokens and then performing prediction based on the remaining *rationale* (Lei et al., 2016; Bastings et al., 2019; Jain et al., 2020), or jointly performing prediction and rationale generation in cases where datasets with annotated rationales are available (Camburu et al., 2018). For some time, the attention mechanism (Bahdanau et al., 2014) has also been considered inherently interpretable. However, the jury is still out on whether such explanations can be considered faithful (Jain and Wallace, 2019; Wiegrefe and Pinter, 2019; Tutek and Šnajder, 2020; Bastings and Filippova, 2020).

*Faithfulness* is one of the most important desiderata of explanation methods (Jacovi and Goldberg, 2020) – faithful explanations are those that are true to the inner decision-making process of the model. Approaches to evaluating faithfulness rely on measuring how the confidence of the model changes when inputs are perturbed (Kindermans et al., 2019) or completely dropped from the model (Li et al., 2016; Serrano and Smith, 2019). However, perturbations to input often result in corrupted instances that fall off the data manifold and appear nonsensical to humans (Feng et al., 2018) or fail to identify all salient tokens properly (ROAR; Hooker et al., 2019) – raising questions about the validity of perturbation-based evaluation. Recursive ROAR (Madsen et al., 2022) alleviates the issues of its predecessor at the cost of requiring many prohibitively expensive retraining steps, further motivating us to seek efficient solutions which do not require retraining the model multiple times. Another option is to leverage the *evaluation-by-agreement* (Jain and Wallace, 2019) paradigm, which states that an interpretability method should be highly agreeable with other methods to be considered faithful. However, since empirical evidence has shown that

saliency methods exhibit poor agreement between their explanations (Neely et al., 2021), Atanasova et al. (2020) recommend practitioners consider alternative methods for evaluating the quality of interpretability methods, such as diagnostic tests. Finally, methods such as data staining (Sippy et al., 2020) and lexical shortcuts (Bastings et al., 2022) artificially introduce tokens that act as triggers for certain classes – creating a ground truth for faithfulness which can be used as a comparison. Nevertheless, such methods have a certain drawback in that they only offer the ground truth importance of a few artificially inserted tokens, but offer no insight regarding the relative importance of the remainder of the input. Each of the aforementioned methods for estimating faithfulness of interpretability methods has its drawbacks (Jacovi and Goldberg, 2020), and we argue each should be taken in conjunction with others to increase the credibility of their collective verdict.

### 3 Preliminaries

In this section, we delineate our experimental setup, detailing the considered datasets, models, their training procedure, the saliency methods which we use to interpret the decisions of the models, and the regularization techniques we use to improve agreement between saliency methods.

#### 3.1 Datasets

Leaning on the work of Neely et al. (2021), which motivated us to explore the valley of explainability, we aim to investigate the protruding problem of low agreement between saliency methods. We investigate three different types of single-sequence binary classification tasks on a total of four datasets. In particular, we evaluate sentiment classification on the movie reviews (IMDB; Maas et al., 2011) and the Stanford Sentiment Treebank (SST-2; Socher et al., 2013) datasets, using the same data splits as Jain and Wallace (2019). We include two more tasks, examining the subjectivity dataset (SUBJ; Pang and Lee, 2004), which classifies movie snippets into subjective or objective, and question type classification (TREC; Li and Roth, 2002). To frame the TREC task as binary classification, we select only the examples labeled with the two most frequent classes (ENTY – entities, HUM – human beings) and discard the rest.

#### 3.2 Models

For comparability, we opt for the same models as Neely et al. (2021). Specifically, we employ the Bi-LSTM with additive self-attention (JWA; Jain and Wallace, 2019). We initialize word representations for the JWA model to 300-d GloVe embeddings (Pennington et al., 2014). We also employ a representative model from the Transformer family (Vaswani et al., 2017) in DistilBERT (DBERT; Sanh et al., 2019).

Both models work similarly: the input sequence of tokens  $\{x_1, \dots, x_T\}$  is first embedded  $\{e_1, \dots, e_T\}$  and then contextualized  $\{h_1, \dots, h_T\}$  by virtue of an LSTM network or a Transformer. The sequence of contextualized hidden states is then aggregated to a sequence representation  $h$ , which is then fed as input to a decoder network.

#### 3.3 Explainability Methods

We make use of ready-made explainability methods from the propagation- and gradient-based families used by Neely et al. (2021): Deep-LIFT (Shrikumar et al., 2017), Integrated Gradients (Int-Grad; Sundararajan et al., 2017) and their Shapley variants (Lundberg and Lee, 2017), Deep-SHAP and Grad-SHAP.<sup>2</sup> Since we evaluate agreement on the entire test set instead of an instance subset (Neely et al., 2021), we exclude LIME (Ribeiro et al., 2016) from the comparison as it is not computationally feasible to train the surrogate model for all test instances across all training setups.

Each saliency method produces a set of importance scores for each input (sub)word token. When evaluating the agreement between different saliency methods for a single trained model, one would expect the importance scores for the same input instance to be similar, as the same set of parameters should produce a unique and consistent importance ranking of input tokens.

#### 3.4 Regularization Methods

As alluded to earlier, we suspect one cause of disagreement between saliency method explanations to be rooted in representation entanglement. To counteract this issue, we employ two regularization schemes that have been shown to improve the faithfulness of the attention mechanism as a method

<sup>2</sup>We use implementations of explainability methods from the Captum framework: <https://github.com/pytorch/captum>

of interpretability: CONICITY (Mohankumar et al., 2020) and TYING (Tutek and Šnajder, 2020). Both of these methods address what we believe is the same underlying issue in *recurrent* models – the fact that hidden representations  $h_t$  are often very similar to each other, indicating that they act more as a sequence representation rather than a contextualization of the corresponding input token  $x_t$ .

Each regularization method tackles this problem in a different manner. CONICITY aims to increase the angle between each hidden representation and the mean of the hidden representations of a single instance. The authors first define the *alignment to mean* (ATM) for each hidden representation as the cosine similarity of that representation to the average representation:

$$\text{ATM}(h_i, \mathbf{H}) = \text{cosine}\left(h_i, \frac{1}{T} \sum_{j=1}^T h_j\right) \quad (1)$$

where  $\mathbf{H} = \{h_1, \dots, h_T\}$  is the set of hidden representations for an instance of length  $T$ . Conicity is then defined as the average ATM for all hidden states  $h_i \in H$ :

$$\text{conicity}(\mathbf{H}) = \frac{1}{T} \sum_{i=1}^T \text{ATM}(h_i, H) \quad (2)$$

A conicity value implies that all hidden representations exist in a narrow cone and have high similarity – to counteract this unwanted effect, during training, we minimize this regularization term weighted by  $\lambda_{con}$  along with the binary cross entropy loss.

Similarly, TYING also aims to incentivize differences between hidden states by enforcing them to “stay true to their word” through minimizing the  $L_2$  norm of the difference between each hidden state  $h_t$  and the corresponding input embedding  $e_t = \text{embed}(x_t)$ :

$$\text{tying}(\mathbf{H}, \mathbf{E}) = \frac{1}{T} \sum_{i=1}^T \|h_i - e_i\|_2^2 \quad (3)$$

where  $\mathbf{E} = \{e_1, \dots, e_T\}$  is the sequence of embedded tokens. During training, we minimize this regularization term weighted by  $\lambda_{tying}$  along with the binary cross entropy loss.

By penalizing the difference between hidden representations and input embedding, one achieves two goals: (1) the embedding and hidden state representation spaces become better aligned, and (2) each hidden representation comes closer to its input

embedding. The latter enforces hidden states to differ from each other: because different embeddings represent the semantics of different tokens, their representations should also differ, and this effect is then also evident in the hidden representations.

Although both works introduced other methods of enforcing differences between hidden states, namely orthogonal-LSTM and masked language modeling as an auxiliary task, we opt for CONICITY and TYING as they were both shown to be more efficient and more stable in practice.

## 4 Improving Agreement

In this section, we present two modifications of the existing *evaluation-by-agreement* procedure: (1) complementing rank-correlation with a linear correlation measure more robust to rank changes caused by small differences in importance weights, and (2) regularizing the models with the goal of reducing entanglement in the hidden space, and as a consequence, improving agreement.

### 4.1 Choice of Correlation Metric

Previous work (Jain and Wallace, 2019; Neely et al., 2021) has evaluated the agreement between two explainability methods by using rank-correlation as measured by Kendall- $\tau$  (Kendall, 1938). Although Kendall- $\tau$  is generally more robust than Spearman’s rank correlation, i.e., it has smaller gross-error sensitivity (Croux and Dehon, 2010), we still face difficulties when using Kendall- $\tau$  for evaluating agreement. As Jain and Wallace (2019) also note, perturbations in ranks assigned to tokens in the tail of the saliency distribution have a large influence on the agreement score. In addition, rankings are also unstable when saliency scores for the most relevant tokens are close to one another. In Figure 1, we illustrate the deficiencies of using rank correlation on a toy example of explaining sentiment classification. While saliency scores attributed to tokens differ slightly, the differences in rank order are significant, lowering agreement according to Kendall- $\tau$  due to the discretization of raw saliency scores when converted into ranks. We believe that a better approach to approximating agreement is to use a linear correlation metric such as Pearson’s  $r$ , as it evaluates whether both saliency methods assign similar importance scores to the same tokens – which is a more robust setup if we assume small amounts of noise in importance attribution between different methods.



		D-SHAP		G-SHAP		Int-Grad	
		$k_\tau$	$p_r$	$k_\tau$	$p_r$	$k_\tau$	$p_r$
DeepLIFT	SUBJ	1.	1.	.31	.45	.43	<b>.64</b>
	SST	1.	1.	.30	.47	.35	<b>.54</b>
	TREC	1.	1.	.12	.31	.15	.33
	IMDB	1.	1.	.29	<b>.59</b>	.28	<b>.60</b>
D-SHAP	SUBJ			.31	.45	.43	<b>.64</b>
	SST			.30	.47	.35	<b>.54</b>
	TREC			.12	.31	.15	.33
	IMDB			.29	<b>.60</b>	.28	<b>.60</b>
G-SHAP	SUBJ					<b>.62</b>	<b>.78</b>
	SST					<b>.70</b>	<b>.87</b>
	TREC					<b>.66</b>	<b>.85</b>
	IMDB					<b>.68</b>	<b>.94</b>

(a) JWA

		D-SHAP		G-SHAP		Int-Grad	
		$k_\tau$	$p_r$	$k_\tau$	$p_r$	$k_\tau$	$p_r$
DeepLIFT	SUBJ	.24	.44	.10	.19	.12	.21
	SST	.19	.34	.09	.17	.10	.20
	TREC	.16	.30	.12	.25	.12	.26
	IMDB	.28	<b>.51</b>	.11	.24	.13	.27
D-SHAP	SUBJ			.11	.22	.13	.24
	SST			.10	.19	.11	.23
	TREC			.13	.28	.14	.30
	IMDB			.12	.26	.14	.30
G-SHAP	SUBJ					.36	<b>.58</b>
	SST					.31	<b>.54</b>
	TREC					.42	<b>.71</b>
	IMDB					.29	<b>.55</b>

(b) DBERT

Table 1: Agreement between pairs of saliency methods in terms of Kendall- $\tau$  ( $k_\tau$ ) and Pearson- $r$  ( $p_r$ ) for BASE variants of (a) JWA and (b) DBERT. We average the agreement over five runs with different seeds. D-SHAP and G-SHAP denote Deep-SHAP and Grad-SHAP, respectively. The values marked in **bold** signify a strong agreement (greater than 0.5), as per individual metrics  $k_\tau$  or  $p_r$ .

We now investigate how Pearson- $r$  ( $p_r$ ) compares to Kendall- $\tau$  ( $k_\tau$ ) when evaluating agreement. In Table 1 we compare agreement scores produced by  $p_r$  and  $k_\tau$  across all datasets for JWA and DBERT, respectively. For JWA, saliency methods display agreement greater than 0.5 only in 8/24 cases according to  $k_\tau$  and in 16/24 cases as per  $p_r$ . For DBERT, these figures are 0/24 and 5/24, respectively. While the overall agreement is subpar, we posit that  $k_\tau$  further exacerbates the evaluation process. We believe this is caused by tokens with approximately equal relative importance assigned slightly different rankings by explainability methods, which  $k_\tau$  harshly penalizes. To address this,

	.01	.02	.03	.04	.05	.06	1.00
A	The	old	Japanese	movie	Rashomon	is	great.
	.06	.05	.04	.03	.02	.01	.80
B	The	old	Japanese	movie	Rashomon	is	great.

Figure 1: A toy example of sentiment classification illustrating the problems with Kendall- $\tau$ . The corresponding agreement on the shown example is  $k_\tau = -.43$  and  $p_r = .99$ . Token opacity indicates a higher saliency score, i.e., token relevance, written at the top of the token box. Each of the two explainability methods, A and B, outputs its saliency scores. The value of Kendall- $\tau$  is much lower than Pearson’s correlation because the irrelevant tokens are perturbed, despite the fact that the tokens are correctly partitioned into more important and less important ones.

we recommend using the Pearson correlation coefficient as an additional measure in the evaluation of agreement, as it is more robust to rank changes caused by small differences in saliency scores.

## 4.2 Regularizing Models

Our next goal is to improve agreement between saliency methods through intervention in the training procedure, namely by applying regularization to promote disentanglement in the hidden space. In Table 2 we report correlation scores on the test splits of all datasets for regularized models (CONICITY, TYING) and their unregularized variants (BASE). We notice that both regularization techniques have a positive effect on agreement across both correlation metrics, indicating that regularization techniques alleviate a deeper issue that also affects the interpretability of attention weights. In Table 3 we report  $F_1$  scores on the test set for the regularized and unregularized models with the best performance on the validation split. We observe that regularized models generally perform comparably well to unregularized ones on downstream tasks, indicating that the improvement in the agreement does not come at a cost for downstream performance. When selecting regularized models, we choose ones with the strongest regularization scale hyperparameter that performs within 3  $F_1$  points on the validation set compared to the unregularized model (cf. details in Appendix A.2).

## 5 The Cartography of Agreement

We have shown that by using a more appropriate correlation measure and applying regularization, the agreement of saliency methods increases sig-

		$\tau_k$			$r_p$		
		B	C	T	B	C	T
JWA	SUBJ	.52	.48	<b>.65</b> <sup>†</sup>	.66	.70	<b>.88</b> <sup>†</sup>
	SST	.50	.67	<b>.68</b>	.65	<b>.90</b> <sup>†</sup>	.86
	TREC	.37	<b>.77</b> <sup>†</sup>	.68	.52	<b>.98</b> <sup>†</sup>	.93
	IMDB	.47	.52	<b>.60</b> <sup>†</sup>	.72	.64	<b>.80</b> <sup>†</sup>
DBERT	SUBJ	.18	.28	<b>.36</b> <sup>†</sup>	.31	.48	<b>.57</b> <sup>†</sup>
	SST	.15	.15	<b>.33</b> <sup>†</sup>	.28	.27	<b>.60</b> <sup>†</sup>
	TREC	.18	.17	<b>.28</b> <sup>†</sup>	.35	.34	<b>.53</b> <sup>†</sup>
	IMDB	.18	.20	<b>.24</b> <sup>†</sup>	.36	.42	<b>.51</b> <sup>†</sup>

Table 2: Average agreement over all pairs of saliency methods. We report agreement in terms of Kendall- $\tau$  ( $k_\tau$ ) and Pearson- $r$  ( $r_p$ ) from the last epoch for three model variants: B – base, C – conicity, and T – tying. **Bold** numbers indicate highest agreement among the three model flavors. Results are averages over 5 runs with different seeds. We ran one-sided Wilcoxon signed-rank tests to check for statistical significance. Agreement values significantly higher ( $p < .05$ ) than the other variants in two individual Wilcoxon tests are marked with a  $\dagger$ . We adjusted the  $p$ -values for family-wise error rate due to multiple tests using the Holm-Bonferroni method.

		Base	Conicity	Tying
DBERT	SUBJ	.93 <sub>.01</sub>	.90 <sub>.02</sub>	.93 <sub>.00</sub>
	SST	.83 <sub>.00</sub>	.83 <sub>.01</sub>	.82 <sub>.01</sub>
	TREC	.92 <sub>.01</sub>	.92 <sub>.01</sub>	.91 <sub>.01</sub>
	IMDB	.86 <sub>.01</sub>	.86 <sub>.01</sub>	.88 <sub>.00</sub>
JWA	SUBJ	.92 <sub>.00</sub>	.90 <sub>.00</sub>	.89 <sub>.00</sub>
	SST	.78 <sub>.04</sub>	.76 <sub>.02</sub>	.78 <sub>.02</sub>
	TREC	.89 <sub>.02</sub>	.86 <sub>.01</sub>	.89 <sub>.01</sub>
	IMDB	.89 <sub>.00</sub>	.88 <sub>.00</sub>	.86 <sub>.00</sub>

Table 3:  $F_1$  scores on *test* sets across datasets for DBERT and JWA. We report the test results on epochs in which the model had the best performance on the validation set. Columns correspond to the base and regularized models. Numbers in subscript denote standard deviation on 5 runs with different seeds.

nificantly. In this section, we are interested in finding out the cause of increased agreement obtained through applying regularization – are there certain instance groups in the dataset that benefit the most, and if so, what changes in the representation space resulted in the increased agreement? We leverage methods from dataset cartography (Swayamdipta et al., 2020) to distribute instances into *easy-to-learn*, *hard-to-learn*, and *ambiguous* categories based on their prediction confidence and variability. Concretely, if an instance exhibits **low** prediction variability and **high** prediction confidence between epochs, this implies that the model can quickly and

accurately classify those instances, making them *easy-to-learn*. Instances that also exhibit **low** variability but **low** prediction confidence, align with the idea that the model is consistently unable to correctly classify them, making them *hard-to-learn*. Finally, instances that exhibit **high** variability and confidence close to the decision threshold indicate that the model is likely often changing its prediction between class labels for those instances, making them *ambiguous*. Since ambiguous instances are characterized by confidence near the prediction threshold, Swayamdipta et al. (2020) complement *variability* and *confidence* with another statistic introduced by Chang et al. (2017), namely *closeness*, defined as  $c_i = p^{(i)} \cdot (1 - p^{(i)})$ , where  $p^{(i)}$  is the average correct class probability of instance  $x^{(i)}$  across all training epochs. A **high** closeness value denotes that the instance is consistently near the decision boundary and, thus, is a good indicator of ambiguity within the model.

Intuitively, one would expect high agreement between saliency methods on instances that are easy to learn and low agreement otherwise. However, we find the converse is true when computing how agreement distributes across instance groups. In unregularized models, we observe that easy-to-learn instances exhibit low average agreement, while ambiguous instances have a high average agreement. In Table 4, we report average agreement scores across all pairs of saliency methods on representative samples from each cartography group.<sup>3</sup> We observe a clear distinction in agreement for both the base and regularized models, which is higher for ambiguous instances when compared to easy- and hard-to-learn instance groups. Furthermore, we can observe a consistently high increase in agreement when the models are regularized across all instance groups for all datasets, indicating that regularization techniques reduce representation entanglement.

One might wonder how the increase in agreement distributes across instances and dataset cartography attributes. In Figure 2, we visualize how the relationship between agreement and cartography attributes changes when the models are regularized. We observe that for the JWA model, all datasets ex-

<sup>3</sup>We select representative samples for each group through the relative frequency of their correct classification. If out of 5 epochs, an instance was correctly classified 5 times, it is representative of the *easy-to-learn* category. If it was correctly classified 0 times, it is representative of the *hard-to-learn* category, and if the number of correct classifications is 2 or 3, it is representative of the *ambiguous* category.

	Easy		Amb		Hard	
	B	T	B	T	B	T
SUBJ	.28	<b>.52</b> <sup>†</sup>	.48	<b>.74</b> <sup>†</sup>	.40	<b>.57</b> <sup>†</sup>
SST	.24	<b>.57</b> <sup>†</sup>	.36	<b>.63</b> <sup>†</sup>	.30	<b>.55</b> <sup>†</sup>
TREC	.33	<b>.49</b> <sup>†</sup>	.50	<b>.62</b> <sup>†</sup>	.32	<b>.40</b> <sup>†</sup>
IMDB	.34	<b>.48</b> <sup>†</sup>	.42	<b>.59</b> <sup>†</sup>	.36	<b>.51</b> <sup>†</sup>

Table 4: Average agreement (Pearson- $r$ ) across saliency methods pairs per cartography groups. We select representative samples for each group based on the number of times a certain instance was classified correctly during training. We report average agreement for the unregularized model (B) and the one regularized by weight tying (T), with the numbers in **bold** indicating the higher agreement value among the two models. We averaged the results over 5 runs. We ran one-sided Wilcoxon tests to check whether T is significantly better than B for a particular group. Significantly higher agreement values ( $p < .05$ ) are marked with a †.

hibit a consistent and significant increase in agreement. Furthermore, we notice that for the DBERT model, apart from increasing the agreement, regularization reduces the confidence of the model predictions and increases variability – indicating that it reduces the known problem of overconfidence present in pre-trained language models.

### 5.1 The Curvature of Agreement

To better understand the cause of this distinction between various feature groups, we now analyze local curvature and density in the representation space. We are interested in: (1) how densely the instances are distributed in the representation space across cartography categories and (2) whether the local space around an instance is sharp or smooth. For both models and all instances, we obtain sequence representations  $h$  used as inputs to the decoder. We estimate instance density as the average distance to the nearest instance in the dataset. We estimate local smoothness around an instance representation as the  $L_2$  norm of the gradient of the hidden representation with respect to the input embeddings. If the gradient norm is high, the local space is sharp and minor perturbations can have a large effect on the prediction probability.

In Table 5, we report correlations between each of these two statistics and dataset cartography attributes. We observe that for the unregularized model, there is a significant negative correlation between confidence and both gradient norm and minimum distance to the nearest example, indicating that the local space around easy instances is smooth

and densely populated. On the other hand, there is a high positive correlation between both closeness and variability and both gradient norm and minimum distance to the nearest example – indicating that the local space around ambiguous instances is sharp and sparsely populated. When we turn our attention to the regularized model, we observe that the correlation between the gradient norm and any of the cartography attributes vanishes, while the correlations between distance and the attributes are reduced in absolute value and their sign is flipped.

From these observations, we hypothesize that the cause of low agreement on easy-to-learn instances is the multitude of possible explanations as to why such an instance should be correctly classified. From the viewpoint of *plausibility*, this hypothesis is in line with the Rashomon effect (Breiman, 2001) which is about there often existing a multitude of adequate descriptions that end up with the same error rate, or in our case, prediction probability – however it should not apply to *faithfulness*, as a single model instance should adhere to a single explanation. However, due to a plethora of corroborating evidence for easy-to-learn instances, the representation space around them is smooth to such an extent that perturbations do not significantly affect the prediction probability, which in turn adversely affects gradient- and propagation-based explanation methods. The converse is true for ambiguous instances, where we hypothesize the model observes evidence for both classes and is unable to reach a confident decision. However, this difficulty in reaching a decision also causes saliency methods to have a precise definition of what the evidence is – as the local curvature is sharp, and any minor perturbation could significantly affect prediction probability. We believe that local curvature statistics could be used as a metric for measuring whether a trained model is better suited to analysis through explainability methods.

## 6 Conclusion

We analyzed two prototypical models from different families in JWA and DBERT with the goal of finding out the cause of low agreement between saliency method interpretations. We first take a closer look at Kendall- $\tau$ , the previously used rank-order correlation metric, and demonstrate that it can be prone to exhibiting high differences in agreement for small perturbations in importance scores. To account for this, when analyzing agreement be-

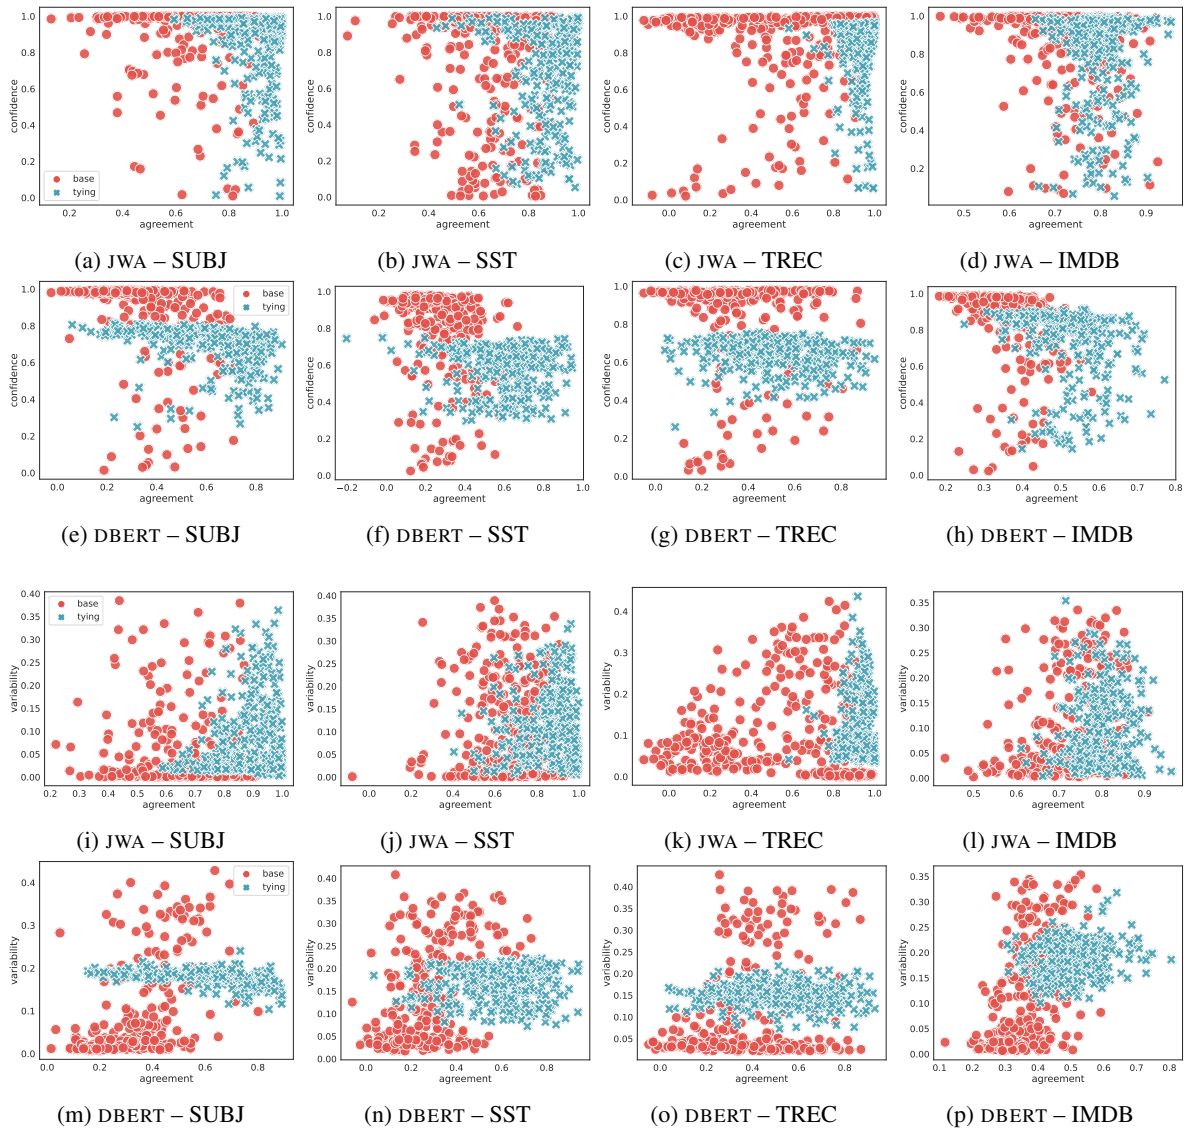


Figure 2: Relationship between agreement and cartography. Subfigures (a)–(h) and (i)–(p) show Pearson- $r$  agreement for each instance in the dataset with respect to *confidence* and *variability*, respectively. The red dots pertain to the unregularized model (BASE), and the blue crosses to the regularized model (TYING).

	Conf		Close		Var	
	B	T	B	T	B	T
Grad norm	-.39	-.02	.52	.05	.46	.06
Min dist	-.53	.25	.72	-.39	.64	.16

Table 5: Correlations (Pearson- $r$ ) between local curvature statistics in the representation space and cartography attributes. We report average gradient norms (*grad norm*) of the hidden representation with respect to the input embeddings and average distance to the nearest instance (*min dist*). Columns correspond to the cartography attributes: *conf* – confidence, *close* – closeness, and *var* – confidence variance. We report the results for the BASE model (B) and the model regularized by TYING (T). Results reported are averages over all datasets.

tween saliency methods, we propose researchers include a linear correlation metric such as Pearson- $r$ , which is robust to small importance perturbations. Taking a step further, we applied two regularization techniques, TYING and CONICITY, originally aimed at increasing faithfulness of attention explanations, with the hypothesis that the issue underpinning disagreements and unfaithfulness is the same – representation entanglement in the hidden space. We showed that regularization consistently and significantly improves agreement scores across all models and datasets with a minimal penalty for classification performance. Having demonstrated that it is possible to improve upon the low agreement scores, we attempted to offer intuition on



which instance categories saliency methods agree the least and show that surprisingly, *easy-to-learn* instances are *hard to agree* on. Lastly, we offered insights into how the representation space morphs when regularization is applied and linked these findings with dataset cartography categories, paving the way for further work on understanding what properties of neural models affect interpretability.

## Limitations

Our work has a number of important limitations that affect the conclusions we can draw from it. First and foremost, evaluating the faithfulness of model interpretations is problematic as we do not have ground truth annotations for token importances. Thus, when applying the *agreement-as-evaluation* paradigm, we implicitly assume that most saliency methods are close to the truth – an assumption that we cannot verify. However, every method of evaluating faithfulness has its own downsides. Token and representation erasure runs the risk of drawing conclusions from corrupted inputs that fall off the data manifold. We argue that while *agreement-as-evaluation* is far from an ideal way of evaluating faithfulness, it still increases credibility when used along with other techniques.

Secondly, our work is limited both with respect to the datasets and models considered. Specifically, we only evaluate one Transformer-based model from the masked language modeling family, and it is entirely possible that the findings do not generalize to models pre-trained on different tasks. Also, we only consider single sequence classification datasets – mainly due to the fact that the issues with the faithfulness of attention were most prevalent in those setups, which we assumed would be the same for agreement due to the same hypothesized underlying issue. We believe that tasks that require retention of token-level information in hidden states, such as sequence labeling and machine translation, would exhibit higher agreement overall, even without intervention through regularization. We leave this analysis for future work.

## References

Samira Abnar and Willem Zuidema. 2020. Quantifying attention flow in transformers. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4190–4197.

Pepa Atanasova, Jakob Grue Simonsen, Christina Lioma, and Isabelle Augenstein. 2020. 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. Association for Computational Linguistics.

Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. 2015. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 10(7):e0130140.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.

Jasmijn Bastings, Wilker Aziz, and Ivan Titov. 2019. Interpretable neural predictions with differentiable binary variables. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2963–2977, Florence, Italy. Association for Computational Linguistics.

Jasmijn Bastings, Sebastian Ebert, Polina Zablotskaia, Anders Sandholm, and Katja Filippova. 2022. “will you find these shortcuts?” a protocol for evaluating the faithfulness of input salience methods for text classification. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 976–991, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Jasmijn Bastings and Katja Filippova. 2020. 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.

Leo Breiman. 2001. Statistical modeling: The two cultures (with comments and a rejoinder by the author). *Statistical science*, 16(3):199–231.

Oana-Maria Camburu, Tim Rocktäschel, Thomas Lukasiewicz, and Phil Blunsom. 2018. e-snli: Natural language inference with natural language explanations. *Advances in Neural Information Processing Systems*, 31.

Rich Caruana, Yin Lou, Johannes Gehrke, Paul Koch, Marc Sturm, and Noemie Elhadad. 2015. Intelligent 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*, pages 1721–1730.

Haw-Shiuan Chang, Erik Learned-Miller, and Andrew McCallum. 2017. Active bias: Training more accurate neural networks by emphasizing high variance samples. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

- Christophe Croux and Catherine Dehon. 2010. Influence functions of the spearman and kendall correlation measures. *Statistical methods & applications*, 19(4):497–515.
- Misha Denil, Alban Demiraj, and Nando De Freitas. 2014. Extraction of salient sentences from labelled documents. *arXiv preprint arXiv:1412.6815*.
- Shi Feng, Eric Wallace, Alvin Grissom II, Mohit Iyyer, Pedro Rodriguez, and Jordan Boyd-Graber. 2018. Pathologies of neural models make interpretations difficult. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3719–3728.
- Rory Mc Grath, Luca Costabello, Chan Le Van, Paul Sweeney, Farbod Kamiab, Zhao Shen, and Freddy Lecue. 2018. Interpretable credit application predictions with counterfactual explanations. *arXiv preprint arXiv:1811.05245*.
- Sara Hooker, Dumitru Erhan, Pieter-Jan Kindermans, and Been Kim. 2019. A benchmark for interpretability methods in deep neural networks. *Advances in neural information processing systems*, 32.
- Alon Jacovi and Yoav Goldberg. 2020. **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. Association for Computational Linguistics.
- 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)*, pages 3543–3556.
- Sarthak Jain, Sarah Wiegrefe, Yuval Pinter, and Byron C Wallace. 2020. Learning to faithfully rationalize by construction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4459–4473.
- Danielle Leah Kehl and Samuel Ari Kessler. 2017. Algorithms in the criminal justice system: Assessing the use of risk assessments in sentencing.
- Maurice G Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.
- Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. 2019. The (un) reliability of saliency methods. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pages 267–280. Springer.
- Diederik P. Kingma and Jimmy Ba. 2015. **Adam: A method for stochastic optimization**. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Tao Lei, Regina Barzilay, and Tommi Jaakkola. 2016. **Rationalizing neural predictions**. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 107–117, Austin, Texas. Association for Computational Linguistics.
- Jiwei Li, Will Monroe, and Dan Jurafsky. 2016. Understanding neural networks through representation erasure. *arXiv preprint arXiv:1612.08220*.
- Xin Li and Dan Roth. 2002. **Learning question classifiers**. In *COLING 2002: The 19th International Conference on Computational Linguistics*.
- Ilya Loshchilov and Frank Hutter. 2017. **Fixing weight decay regularization in adam**. *CoRR*, abs/1711.05101.
- Scott M Lundberg and Su-In Lee. 2017. **A unified approach to interpreting model predictions**. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. 2011. **Learning word vectors for sentiment analysis**. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA. Association for Computational Linguistics.
- Andreas Madsen, Nicholas Meade, Vaibhav Adlakha, and Siva Reddy. 2022. **Evaluating the faithfulness of importance measures in NLP by recursively masking allegedly important tokens and retraining**. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1731–1751, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Ana Marasovic, Iz Beltagy, Doug Downey, and Matthew Peters. 2022. **Few-shot self-rationalization with natural language prompts**. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 410–424, Seattle, United States. Association for Computational Linguistics.
- Clara Meister, Stefan Lazov, Isabelle Augenstein, and Ryan Cotterell. 2021. 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.
- Akash Kumar Mohankumar, Preksha Nema, Sharan Narasimhan, Mitesh M Khapra, Balaji Vasan Srinivasan, and Balaraman Ravindran. 2020. Towards transparent and explainable attention models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4206–4216.

- Michael Neely, Stefan F Schouten, Maurits JR Bleeker, and Ana Lucic. 2021. Order in the court: Explainable ai methods prone to disagreement. *arXiv preprint arXiv:2105.03287*.
- Bo Pang and Lillian Lee. 2004. [A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts](#). In *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*, pages 271–278, Barcelona, Spain.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. [Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter](#). *CoRR*, abs/1910.01108.
- Sofia Serrano and Noah A Smith. 2019. Is attention interpretable? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2931–2951.
- Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. 2017. Learning important features through propagating activation differences. In *International conference on machine learning*, pages 3145–3153. PMLR.
- Jacob Sippy, Gagan Bansal, and Daniel S Weld. 2020. Data staining: A method for comparing faithfulness of explainers. In *Proceedings of ICML Workshop on Human Interpretability in Machine Learning*.
- Richard Socher, John Bauer, Christopher D. Manning, and Andrew Y. Ng. 2013. [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. Association for Computational Linguistics.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *International Conference on Machine Learning*, pages 3319–3328. PMLR.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A. Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9275–9293, Online. Association for Computational Linguistics.
- Martin Tutek and Jan Šnajder. 2020. Staying true to your word:(how) can attention become explanation? In *Proceedings of the 5th Workshop on Representation Learning for NLP*, pages 131–142.
- 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*, volume 30. Curran Associates, Inc.
- Sarah Wiegrefe and Yuval Pinter. 2019. 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.

		Base	Conicity	Tying
DBERT	SUBJ	.94	.89	.94
	SST	.85	.85	.85
	TREC	.94	.89	.91
	IMDB	.90	.89	.89
JWA	SUBJ	.93	.90	.91
	SST	.82	.79	.81
	TREC	.91	.87	.89
	IMDB	.90	.87	.88

Table 6:  $F_1$  scores on *validation* sets across datasets for DBERT and JWA. We average the results over 5 runs with different seeds. The scores pertain to the same experiments as in Table 3, where we report test  $F_1$  scores.

## A Reproducibility

### A.1 Experimental Results

#### A.1.1 Setup

For both JWA and DBERT, we use the same pre-processing pipeline on all four datasets. First, we filter out instances with fewer than three tokens to achieve stable agreement evaluation.<sup>4</sup> Next, we lowercase the tokens, remove non-alphanumeric tokens, and truncate the sequence to 200 tokens if the sequence length exceeds this threshold. We set the maximum vocabulary size to 20k for models which do not leverage subword vocabularies.

#### A.1.2 Validation set performance

We report the validation set performance in Table 6.

#### A.1.3 Computing infrastructure

We conducted our experiments on  $2 \times$  AMD Ryzen Threadripper 3970X 32-Core Processors and  $2 \times$  NVIDIA GeForce RTX 3090 GPUs with 24GB of RAM. We used PyTorch version 1.9.0 and CUDA 11.4.

#### A.1.4 Average runtime

Table 9 shows the average experiment runtime for each model across the datasets we used.

#### A.1.5 Number of parameters

The JWA and DBERT models that we used contained 1, 714, 951 and 66, 954, 241 trainable parameters, respectively.

<sup>4</sup>If a sequence consists of only two tokens, rank-correlation with Kendall- $\tau$  will either result in a perfect match, or completely different observations as swapping the two ranks leads to an inverse ranking.

	JWA	DBERT
SUBJ	3.4	11.2
SST	2.7	8.9
TREC	1.2	3.7
IMDB	6.1	107.5

Table 7: Experiment duration in minutes for both models across datasets. We report the average runtime over 5 different runs.

	JWA		DBERT	
	C	T	C	T
SUBJ	1.	1.	5.	1.
SST	1.	0.5	0.1	0.5
TREC	1.	1.	0.1	0.3
IMDB	0.3	1.	1.	1.

Table 8: Selected hyperparameter values for CONICITY (C) and TYING (T).

### A.2 Hyperparameter search

We used the following parameter grids for JWA:  $[10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}, 10^{-5}, 10^{-6}]$  for learning rate, and  $[50, 100, 150, 200]$  for the hidden state dimension. We yield best average results on validation sets across all datasets when the learning rate is set to  $10^{-3}$  and the hidden size is set to 150. For DBERT, we find that the most robust initial learning rate on the four datasets is  $2 \times 10^{-5}$ , among the options we explored  $[5 \times 10^{-4}, 10^{-4}, 10^{-5}, 2 \times 10^{-5}, 5 \times 10^{-5}, 10^{-6}]$ . Additionally, we clip the gradients for both models such that the gradient norm  $\leq 1$ . We use the Adam (Kingma and Ba, 2015) optimizer for JWA and AdamW (Loshchilov and Hutter, 2017) for DBERT. We run both models for 5 epochs and repeat the experiments 5 times with different seeds:  $[1, 2, 3, 4, 5]$ .

For regularization methods, we conducted a grid search with parameter grid  $[0.1, 0.3, 0.5, 1, 5, 10]$  for CONICITY and  $[0.1, 0.3, 0.5, 1, 5, 10, 20]$  for TYING. We select the models with the strongest regularization scale, which is within 3  $F_1$  points from the unregularized model. Table 8 shows the selected values for each model across all datasets.

	JWA	DBERT
SUBJ	3.4	11.2
SST	2.7	8.9
TREC	1.2	3.7
IMDB	6.1	107.5

Table 9: Experiment duration in minutes for both models across datasets. We report the average runtime over 5 different runs.



	Train	Validation	Test	Total
SUBJ	7,000	1,000	2,000	10,000
SST	6,819	868	1,810	9,497
TREC	1,987	159	486	2,632
IMDB	17,212	4,304	4,363	25,879

Table 10: Number of instances in each split and the total number of instances in each dataset after we excluded too short examples (see section 3.1).

### A.3 Dataset statistics

We report the number of instances per split for each dataset in Table 10. We note that all of the datasets we used contain predominantly texts in English.

## B Additional Experiments

We show the full version of local curvature statistics in Table 11 (without averaging over datasets). In Figures 3 to 10 we plot correlation scores ( $k_T$  and  $p_r$ ) with standard deviation on the test splits. We include the results for all datasets across training epochs for regularized models (CONICITY, TYING) when compared to their unregularized, BASE variants.

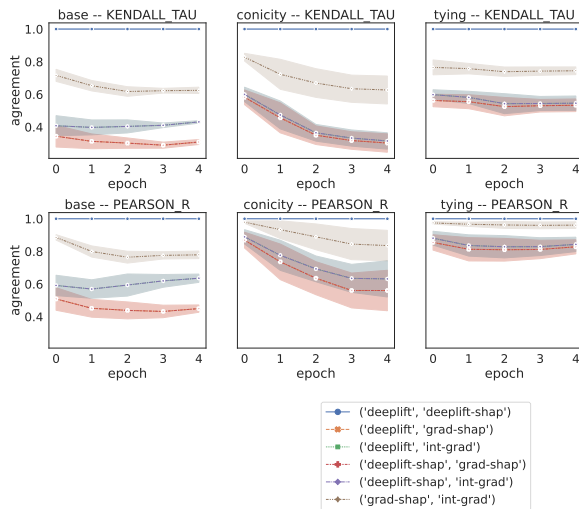


Figure 3: JWA – SUBJ

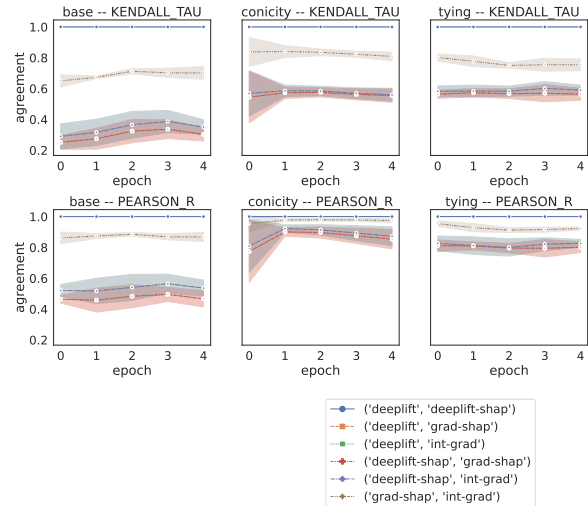


Figure 4: JWA – SST

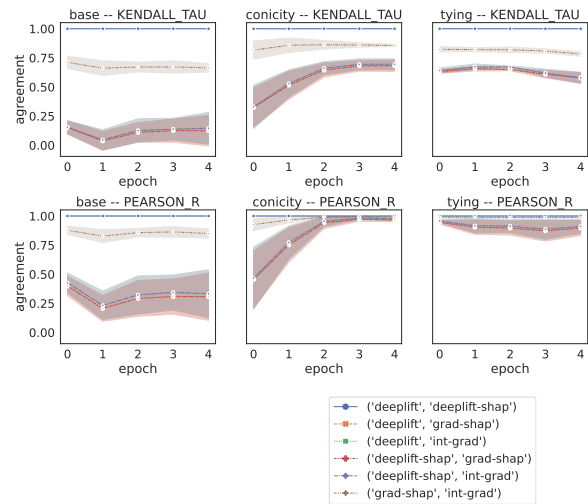


Figure 5: JWA – TREC

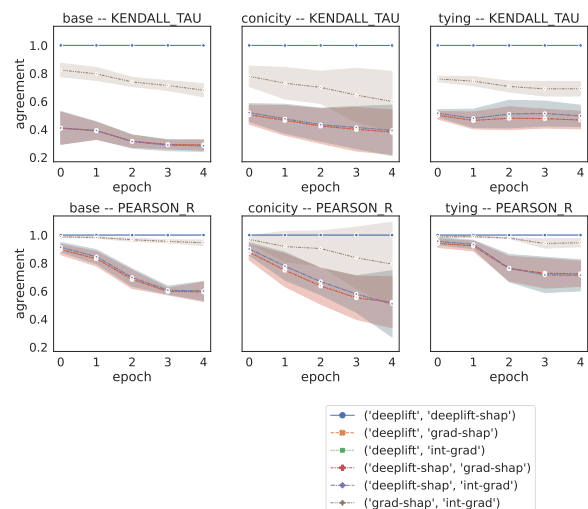


Figure 6: JWA – IMDB

		Confidence		Ambiguity		Variability	
		B	T	B	T	B	T
Grad norm	SUBJ	-.37 <sub>.00</sub>	-.06 <sub>.00</sub>	.48 <sub>.00</sub>	.05 <sub>.02</sub>	.45 <sub>.00</sub>	.00 <sub>.83</sub>
	SST	-.40 <sub>.00</sub>	.02 <sub>.34</sub>	.58 <sub>.00</sub>	.08 <sub>.00</sub>	.42 <sub>.00</sub>	-.23 <sub>.00</sub>
	TREC	-.32 <sub>.00</sub>	.05 <sub>.32</sub>	.39 <sub>.00</sub>	-.12 <sub>.01</sub>	.38 <sub>.00</sub>	.18 <sub>.00</sub>
	IMDB	-.46 <sub>.00</sub>	-.11 <sub>.00</sub>	.61 <sub>.00</sub>	.17 <sub>.00</sub>	.60 <sub>.00</sub>	.30 <sub>.00</sub>
Min dist	SUBJ	-.59 <sub>.00</sub>	.01 <sub>.70</sub>	.79 <sub>.00</sub>	-.10 <sub>.00</sub>	.74 <sub>.00</sub>	.06 <sub>.00</sub>
	SST	-.45 <sub>.00</sub>	.30 <sub>.00</sub>	.70 <sub>.00</sub>	-.44 <sub>.00</sub>	.49 <sub>.00</sub>	.30 <sub>.00</sub>
	TREC	-.55 <sub>.00</sub>	.17 <sub>.00</sub>	.70 <sub>.00</sub>	-.26 <sub>.00</sub>	.68 <sub>.00</sub>	.25 <sub>.00</sub>
	IMDB	-.53 <sub>.00</sub>	.50 <sub>.00</sub>	.70 <sub>.00</sub>	-.74 <sub>.00</sub>	.64 <sub>.00</sub>	.04 <sub>.00</sub>

Table 11: Correlations between local curvature statistics in the representation space and cartography attributes for each dataset. We use average gradient norms (*grad norm*) of the hidden representation with respect to the input embeddings and average distance to the nearest instance (*min dist*). The columns correspond to the cartography attributes. We report the results for the unregularized model (B) and the regularized one to which we applied tying (T). The values in the subscript denote the standard deviation.

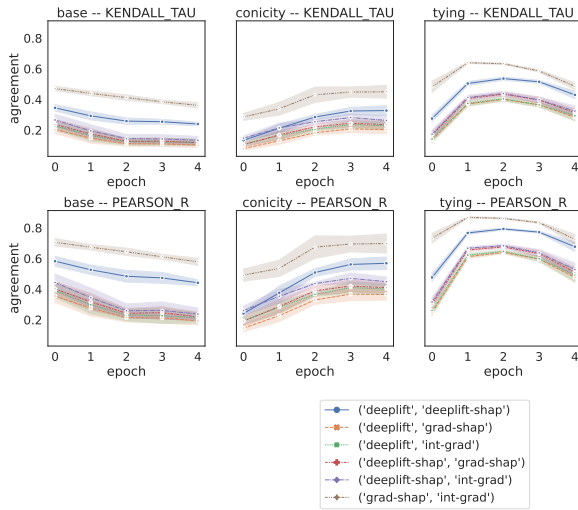


Figure 7: DBERT – SUBJ

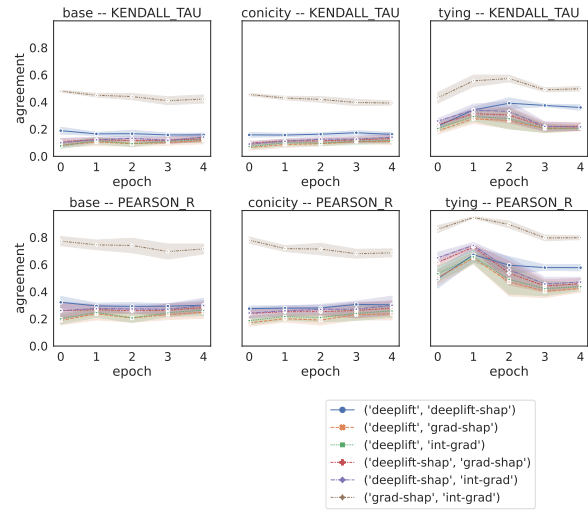


Figure 9: DBERT – TREC

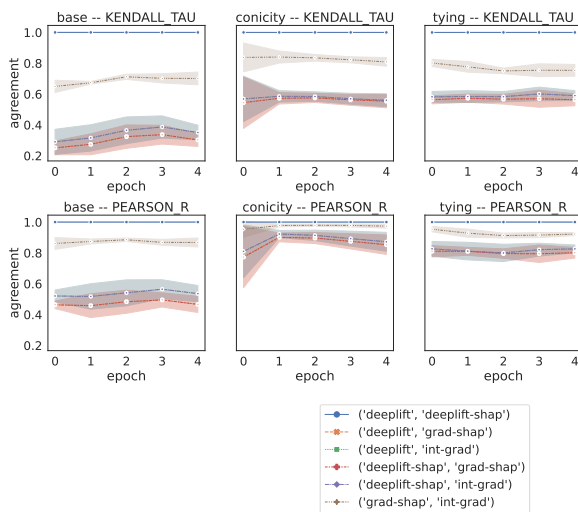


Figure 8: DBERT – SST

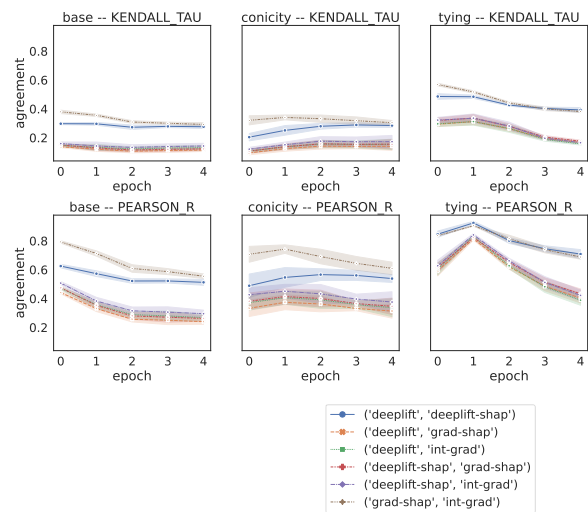


Figure 10: DBERT – IMDB

## ACL 2023 Responsible NLP Checklist

---

### A For every submission:

- A1. Did you describe the limitations of your work?  
*Left blank.*
- A2. Did you discuss any potential risks of your work?  
*Left blank.*
- A3. Do the abstract and introduction summarize the paper’s main claims?  
*Left blank.*
- A4. Have you used AI writing assistants when working on this paper?  
*Left blank.*

### B Did you use or create scientific artifacts?

*Left blank.*

- B1. Did you cite the creators of artifacts you used?  
*Left blank.*
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?  
*Left blank.*
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?  
*Left blank.*
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?  
*Left blank.*
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?  
*Left blank.*
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.  
*Left blank.*

### C Did you run computational experiments?

*Left blank.*

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?  
*Left blank.*

*The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.*

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

*Left blank.*

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

*Left blank.*

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

*Left blank.*

**D**  **Did you use human annotators (e.g., crowdworkers) or research with human participants?**

*Left blank.*

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

*Left blank.*

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

*Left blank.*

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

*Left blank.*

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

*Left blank.*

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

*Left blank.*