

# Assessing Out-of-Domain Language Model Performance from Few Examples

Prasann Singhal\*, Jarad Forristal\*, Xi Ye, and Greg Durrett

Department of Computer Science

The University of Texas at Austin

{prasanns, jarad, xiye, gdurrett}@cs.utexas.edu

## Abstract

While pretrained language models have exhibited impressive generalization capabilities, they still behave unpredictably under certain domain shifts. In particular, a model may learn a reasoning process on in-domain training data that does not hold for out-of-domain test data. We address the task of predicting out-of-domain (OOD) performance in a few-shot fashion: given a few target-domain examples and a set of models with similar training performance, can we understand how these models will perform on OOD test data? We start from the baseline of looking at model accuracy on the few-shot examples, then investigate how to incorporate analysis of the models' behavior using feature attributions to improve our understanding of generalization. Specifically, we explore a set of "factors" designed to reveal model agreement with certain pathological heuristics that may indicate worse generalization capabilities. On textual entailment, paraphrase recognition, and a synthetic classification task, we show that attribution-based factors can help rank relative model OOD performance. However, accuracy on a few-shot test set is a surprisingly strong baseline, particularly when the system designer does not have in-depth prior knowledge about the domain shift.

## 1 Introduction

The question of whether models have learned the right behavior on a training set is crucial for generalization. Deep models have a propensity to learn shallow reasoning shortcuts (Geirhos et al., 2020) like single-word correlations (Gardner et al., 2021) or predictions based on partial inputs (Poliak et al., 2018), particularly for problems like natural language inference (Gururangan et al., 2018; McCoy et al., 2019) and question answering (Jia and Liang, 2017; Chen and Durrett, 2019). Unless we use eval-

\*Equal contribution




Figure 1: Our setting: a system developer is trying to evaluate a collection of trained models on a small amount of hand-labeled data to assess which one may work best in this new domain. Can baselines / attributions help?

uation sets tailored to these spurious signals, accurately understanding if a model is learning them remains a hard problem (Bastings et al., 2021; Kim et al., 2021; Hupkes et al., 2022).

This paper addresses the problem of *predicting* whether a model will work well in a target domain *given only a few examples from that domain*. This setting is realistic: a system designer can typically hand-label a few examples to serve as a test set, but labeling more may be burdensome. Computing accuracy on this small set and using that as a proxy for full-test set performance is a simple baseline for our task, but has high variance, which may cause us to incorrectly rank two models that achieve somewhat similar performance. We hypothesize that we can do better if we can interpret the model's behavior beyond accuracy. With the rise of techniques to analyze post-hoc feature importance in machine-learned models (Lundberg and Lee, 2017; Ribeiro et al., 2016; Sundararajan et al., 2017), we have seen not just better interpretation of models, but improvements such as constraining them to avoid using certain features (Ross et al., 2017) like those associated with biases (Liu and Avci, 2019; Kennedy et al., 2020), or trying to more generally teach the right reasoning process for a problem (Yao et al., 2021; Tang et al., 2021; Pruthi et al.,

2022). If post-hoc interpretation can strengthen a models’ ability to generalize, can they also help us understand it?

Figure 1 illustrates the role this understanding can play. We have three trained models and are trying to rank them for suitability on a new domain. The small labeled dataset is a useful (albeit noisy) indicator of success. However, by checking model attributions on our few OOD samples, we can more deeply understand model behavior and analyze if they use certain pathological heuristics. Unlike past work (Adebayo et al., 2022), we seek to automate this process as much as possible, provided the unwanted behaviors are characterizable by describable heuristics. We use scalar factors, which are simple functions of model attributions, to estimate proximity to these heuristics, similar to characterizing behavior in past work (Ye et al., 2021a). We then evaluate whether these factors allow us to correctly rank the models’ performance on OOD data.

Both on synthetic (Warstadt et al., 2020), and real datasets (McCoy et al., 2019; Zhang et al., 2019), we find that, between models with similar architectures but different training processes, both our accuracy baseline and attribution-based factors are good at distinguishing relative model performance on OOD data. However, on models with different base architectures, we discovering interesting patterns, where factors can very strongly distinguish between different types of models, but cannot always map these differences to correct predictions of OOD performance. In practice, we find probe set accuracy to be a quick and reliable tool for understanding OOD performance, whereas factors are capable of more fine-grained distinctions in certain situations.

**Our Contributions:** (1) We benchmark, in several settings, methods for predicting and understanding relative OOD performance with few-shot OOD samples. (2) We establish a ranking-based evaluation framework for systems in our problem setting. (3) We analyze patterns in how accuracy on a few-shot set and factors derived from token attributions distinguish models.

## 2 Motivating Example

To expand on Figure 1, Figure 2 shows an in-depth motivating example of our process. We show three feature attributions from three different models on an example from the HANS dataset (McCoy et al.,




Figure 2: Explanations generated on the same sample for HANS subsequence data models M1, M2, M3 (have ascending OOD performance). The factor (shaded underlines) from knowledge of the OOD allows us to in this example predict the model ranking.

2019). These models have (unknown) varied OOD performance but similar performance on the in-domain MNLI data (Williams et al., 2018). Our task is then to correctly rank these models’ performance on the out-of-domain HANS dataset in a few-shot manner.

We can consider ranking these models via simple metrics like accuracy on the small few-shot dataset, where higher-scoring models are higher-ranked. However, such estimates can be high variance on small datasets. In Figure 2, only M3 predicts non-entailment correctly, and we cannot distinguish the OOD performance of M1 and M2 without additional information.

Thus, we turn to explanations to gain more insight into the models’ underlying behavior. With faithful attributions, we should be able to determine if the model is following simple inaccurate rules called *heuristics* (McCoy et al., 2019). Figure 2 shows the heuristic where a model predicts that the sentence *A* entails *B* if *B* is a subsequence of *A*. Crucially, we can use model attributions to assess model use of this heuristic :we can sum the attribution mass the model places on subsequence tokens. We use the term *factors* to refer to such functions over model attributions.

The use of factors potentially allows for the automation of detection of spurious signals or shortcut learning (Geirhos et al., 2020). While prior work has shown that spurious correlations are hard for a human user to detect from explanations (Adebayo et al., 2022), well-designed factors could automatically analyze model behavior across a number of tasks and detect such failures.

### 3 Attributions to Predict Performance

In this section, we formalize the ideas presented thus far. Token-level attribution methods (a subset of post-hoc explanations) are methods which, given an input sequence of tokens  $\mathbf{x} \stackrel{\text{def}}{=} x_1, x_2, \dots, x_n$  and a model prediction  $\hat{y} \stackrel{\text{def}}{=} M(\mathbf{x})$  for some task, assign an explanation  $\phi(\mathbf{x}, \hat{y}) \stackrel{\text{def}}{=} a_1, \dots, a_n$  where  $a_i$  corresponds to an attribution or importance score for a corresponding  $x_i$  towards the final prediction. For cases where the model, prediction, and inputs are unambiguous, we abbreviate this simply  $\phi_i \equiv \phi(\mathbf{x}) \stackrel{\text{def}}{=} \phi(\mathbf{x}, M_i(\mathbf{x}))$ .

We assume that the model is trained on an in-domain training dataset  $D_T$  and will be evaluated on some unknown OOD set  $D_O$ . Given two models  $M_0$  and  $M_1$ , with a small amount of data  $D_{(O,t)} \subset D_O$  ( $t = 10$  examples or fewer in our settings), our task is to predict which model will generalize better. We break the process into 2 steps (see Figure 2):

**1. Hypothesize a heuristic.** First we must identify an underlying heuristic  $H$  that reflects pathological model behavior in the OOD dataset. For example, the subsequence heuristic in Figure 2 corresponds to a heuristic which always predicts entailed if the hypothesis is contained within the premise. Let  $h(M_i)$  abstractly reflect how closely the  $i$ th model’s behavior aligns with  $H$ . Let  $s(M_i)$  be the true OOD performance of model  $M_i$ . If we then assume that  $h(M_i)$  faithfully models some pathological heuristic  $H$ , we should have that  $h(M_0) > h(M_1) > \dots > h(M_m)$  implies  $s(M_0) < s(M_1) < \dots < s(M_m)$ . In other words, the more a model  $M_i$  agrees with a pathological heuristic  $H$ , the worse it performs.

**2. Measure alignment.** We now want to predict the ranking of  $s(M_i)$ ; however, with few labeled examples there may be high variance in directly evaluating these metrics. We instead use factors  $f(\mathbf{x}, \phi_i)$  which map tokens and their attributions for model  $M_i$  to scalar scores that should correlate with the heuristic  $H$ . Factors can be designed to align with known pathological heuristics, where higher scores indicate strong model agreement with the associated heuristic. We then estimate the ranking of  $s(M_i)$  using the relative ranking of the corresponding  $h(M_i)$  approximated through factors.

Concretely, to measure the alignment, we first compute for each input  $\mathbf{x}_j \in D_{(O,t)}$  the prediction  $M_i(\mathbf{x}_j)$  and the explanation  $\phi(\mathbf{x}_j)$  for that

prediction. These  $\phi(\mathbf{x}_j)$  are used to compute the score  $f(\mathbf{x}_j, \phi(\mathbf{x}_j))$  for model  $M$ . We take the overall score of the model to be  $F(i) = \frac{1}{t} \sum_{j=1}^t f(\mathbf{x}_j, \phi(\mathbf{x}_j, M_i(\mathbf{x}_j)))$ , the mean over the  $t$  examples in  $D_{(O,t)}$ . We then directly rank models on the basis of the  $F(i)$  values: the higher the average factor value (the more it follows the heuristic), the lower the relative ranking:  $F(0) > F(1) \implies s(M_0) < s(M_1)$ . Therefore we can sort the models by these values and arrive at a predicted ranking. We later also consider factors which do not intuitively map to specific heuristics.

**Baselines** We also consider three principle explanation-agnostic baselines. A natural baseline given  $D_{(O,t)}$  is to simply use the accuracy (ACC) on this dataset:  $\frac{1}{n} \sum_{i=1}^n \mathbb{1}[y_i = M(\mathbf{x}_i)]$ , however this may be noisy on only a few examples, and frequently leads to ties.<sup>1</sup>

We also assess model confidence (CONF), the softmax probability of the predicted label, as well as looking at CONF-GT, the softmax probability of only the ground-truth label.

## 4 Experimental Setup

### 4.1 Models Compared

In this work, we compare various models across different axes yielding different  $D_O$  performance. The first approach we use is **inoculation** (Liu et al., 2019a), which involves fine-tuning models on small amounts or batches of  $D_O$  data alongside in-domain data to increase model performance on OOD data. The second approach we use is varying the model architecture and pre-training (e.g., using a stronger pre-trained Transformer model).

In Section 5, we use inoculation to create 5 RoBERTa-base (Liu et al., 2019b) models of varying  $D_O$  performance for each of the three MSGS sets. In Section 6 where we consider the HANS and PAWS datasets, we inoculate a variety of models. For HANS, we inoculate 5 RoBERTa-large models. We additionally examine DeBERTa-v3-base (He et al., 2021b,a) and ELECTRA-base (Clark et al., 2020) models fine-tuned on in-domain MNLI data. For PAWS, we inoculate 4 RoBERTa-base models on the in-domain  $D_T$  set. We also inoculate ELECTRA-base and DeBERTa-base models. We include complete details for these models in Ap-

<sup>1</sup>Most considered datasets are constructed to mislead models following the heuristic, so this baseline directly measures agreement with a heuristic  $h$ .

pendix A. The generated models represent a realistic problem scenario: a practitioner may have many different models with similar  $D_T$  performance, but different  $D_O$  performance. We specifically crafted suites of models which have both near pairs (models with similar  $D_O$  performance) and far pairs.

## 4.2 Attribution Methods

We experiment using several token-level attribution methods: **LIME** (Ribeiro et al., 2016) computes attribution scores using the weights of a linear model approximating model behavior near a datapoint. **SHAP** (Lundberg and Lee, 2017) is similar to LIME, but uses a procedure using Shapley values. Finally, **Integrated Gradients (TOKIG)** (Sundararajan et al., 2017) compute  $\phi_i$  by performing a line integral over the gradients with respect to token embeddings on a path from a baseline token to the ground truth token; commonly, this baseline token is chosen to be <MASK>. While intuitively sensible, Harbecke (2021) has voiced concerns regarding the use of TOKIG in NLP.

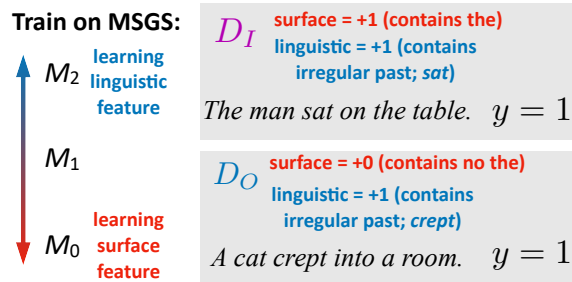
## 4.3 Evaluation Setup

Because model ranking using a small  $D_{(O,t)}$  may be unstable, we conduct all experiments over a number of different sampled  $D_{(O,t)}$  sets. We first sample  $M$  examples from each set (in the range of 200-600), then generate explanations for all models on each example. We then take 400-500 bootstrap samples of size  $n$  (we report results for  $n = 10$ , as experimental results were similar for sizes 5 and 20), simulating many few-shot evaluations. For each bootstrap sample, we analyse  $\binom{m}{2}$  model pairs. Details can be found in Appendix B.

We define a “success” as a technique correctly ranking a model pair, when measured by  $D_O$  performance (on the full set); otherwise is a “failure”. We define *pairwise accuracy* as the accuracy for a method ranking a particular model pair across all bootstrap samples. We define *few-shot accuracy* (or just *accuracy*) as the average of the pairwise accuracies over the  $\binom{m}{2}$  model pairs. By reporting ranking accuracy across a diverse set of models, we ensure a comprehensive evaluation.

## 5 MSGS: A Proof of Concept

We first show experiments on the Mixed Signals Generalization Set (MSGS) dataset presented in Warstadt et al. (2020) as a proof of concept for our methodology. MSGS is a synthetic clas-



- 1) Evaluate on  $D_I$  :  $M_0 = M_1 = M_2$
- 2) Evaluate on  $D_{(O,t)}$  :  $M_0 < M_1 < M_2$  (noisy)
- 3) Attributions on  $D_{(O,t)}$  :  $M_0 < M_1 < M_2$  (correct)

Figure 3: Example from the MSGS train and OOD test sets. The training data conflates a surface and linguistic generalization as described in Warstadt et al. (2020), resulting in models that learn a range of behaviors. Direct evaluation OOD on small data can tell us this, but explanations can also differentiate which of the two patterns is learned and how strongly they are learned.

sification dataset. The training (in-domain) set is composed of sentences where both some linguistic feature (e.g., the presence of an adjective) and a spurious surface feature (e.g., the word “the” being in the sentence) are always associated with a positive label  $y = 1$ . This data is ambiguous, which means the model could rely on either the linguistic or surface feature completely yet still get 100% accuracy on in-domain data. Warstadt et al. (2020) then create sets of OOD data where the linguistic feature becomes associated with the  $y = 1$  positive label, and the surface feature with a  $y = 0$  label. The resulting test accuracy reflects model reliance on one feature or the other. Warstadt et al. (2020) use this to investigate what generalizations are learned at which stages of model pre-training; we investigate whether information from small probe sets can help assess model reliance on the surface feature.

We consider three linguistic features: MORPH (presence of an irregular past verb like “drew”), ADJECT (presence of an adjective), and VERB (if the main verb is an -ing verb), each paired with the surface feature of “the” being in the sentence.

We design factors which look at attributions on the tokens corresponding to these linguistic features, including the tokens surrounding these features as well to account for feature dependence on surrounding words. Our factor  $f(\mathbf{x}, \phi) = -\sum_{i=(m-2)}^{m+2} \phi(x_i)$ , where  $m$  is the index of the feature-critical word for that dataset (e.g., “slept” for IRREG) and  $\phi(x_i)$  is the attribution at an in-

Feature	Method	Accuracy		
MORPH	ACC	90.9		
	CONF	50.9		
	CONF-GT	90.1		
	IRREG	TOKIG 89.2	SHAP 90.6	LIME <b>92.8</b> †
VERB	ACC	94.5		
	CONF	58.0		
	CONF-GT	93.3		
	VERB	TOKIG 92.1	SHAP 94.0	LIME <b>94.9</b>
ADJECT	ACC	89.9		
	CONF	50.5		
	CONF-GT	91.3		
	ADJ	TOKIG 87.4	SHAP 92.1	LIME <b>93.5</b> †

Table 1: Few-shot ranking accuracy metric results on  $D_{(O,t)}$  for MSGS. IRREG, VERB, and ADJ are detailed in Section 5. † indicates statistically significant improvement over accuracy (paired bootstrap test:  $p < 0.05$ )

dex. This factor corresponds closely to the heuristic that the dataset was designed for, or alternately, we can see this factor as inversely proportional to what *other* information the model is using (that is, information outside of this window). We name the factors IRREG, VERB, and ADJ for the MORPH, VERB, ADJECT sets respectively.

Note that this approach assumes that a system designer has prior knowledge of the relevant linguistic and surface feature. This is a generous assumption, and for this dataset is almost sufficient to formulate the rule used to construct it, hence why we call this a proof of concept. We will show more realistic conditions in Section 6.

**Models** To create a suite of models with varying  $D_O$  performance, we inoculate following the steps outlined in Section 4.1. We evaluate our factors via accuracy as described in Section 4.3. More details about the inoculation is present in Section A of the appendix.

**Results** Table 1 shows the results on this dataset. Our ACC baseline performs well: when models differ greatly in performance (e.g., one gets 50% and another gets 90% on the  $D_O$ ), accuracy on the small  $D_{(O,t)}$  ranks these correctly even despite the small subset size. The high regularity of the dataset also means that a model’s behavior does not vary greatly from example to example, further reducing variance. However, this ranking is nevertheless still not perfect. We see that CONF performs very poorly, by contrast, showing that confidence is not helpful

for measuring model behavior.

Overall, we see that methods using explanations are able to beat the ACC baseline, with the exception of TOKIG. We additionally found trends within the explanation techniques themselves, with LIME reliably performing the best, and TOKIG being the worst. But generally, all techniques can offer relevant information, and in the best case, the attributions can tell us more reliably what a model is learning than evaluation on a small set of  $D_{(O,t)}$  data can. In Section 6, we investigate if these results generalize to real-world datasets.

## 6 Realistic OOD Settings

We now consider two datasets corresponding to realistic OOD settings treated in past work.

First, **HANS** (McCoy et al., 2019) targets spurious heuristics within MNLI (Williams et al., 2018), such as the hypothesis being a subsequence of the premise, with balanced test sets that can be used to detect model reliance on these heuristics. Models following these heuristics always predict *entailed* for the hypotheses, and will perform at random chance accuracy on the dataset. We use MNLI as our in-domain training set in this setting.

Second, **PAWS** (Zhang et al., 2019) is a paraphrase identification task. PAWS-QQP is an OOD dataset for Quora Question Pairs (QQP) (Iyer et al., 2017) that is composed of pairs with swapped content words/phrases (e.g., *I ran from the Grand Canyon to California* to *I ran from California to the Grand Canyon*). A paraphrase model that relies heavily on lexical overlap will not be sensitive to these changes, and will always predict the label of  $y = 1$  to indicate paraphrase. We use QQP as our in-domain training set in this setting.

Details regarding models used in this section are presented in Section 4.1. From the test sets of the corresponding datasets, we randomly sample 400 examples from PAWS and 600 from HANS-CON and HANS-SUB each for use in bootstrap sampling, as detailed in Section 4.3. Table 9 provides additional dataset information.

### 6.1 Factors

**General Factors** Both HANS and PAWS involve comparing two sequences  $\mathbf{a}, \mathbf{b}$  of tokens, unlike MSGS which is classification over a single sequence. We can define our input  $\mathbf{x} = a_1, a_2, \dots, a_n, b_1, b_2, \dots, b_m$  as composed of these two sequences  $\mathbf{a}$  and  $\mathbf{b}$  with respective attributions

$\phi_a, \phi_b$ . We evaluate a number of factors that generally target sensitivity to both sequences and their differences, which represent a broad class of potential heuristics.

MAX-DIFF: The difference between maximum attribution in **a** and **b**, i.e.  $\max(\phi_a) - \max(\phi_b)$ .

SUM-DIFF: the difference of the sum of attributions, i.e.  $\sum_{i=1}^n \phi_{a,i} - \sum_{i=1}^m \phi_{b,i}$ .

INDEX-DIFF: The difference of attributions between shared words in **a** and **b**.

FIRST-TOK: The attribution at the the first <SEP> token.

We explicitly note that this is the exhaustive set of factors we experimented with, not a cherry-picked set, in order to provide a comprehensive view of what does and doesn’t work. We crafted these by manually examining attribution patterns on various datasets rather than trying a large number and keeping the best ones.

**HANS Factors** We look at the “subsequence” heuristic discussed in Section 2 and the constituent heuristic, which assumes that the premise entails all complete subtrees in its parse-tree. For the subsequence OOD set (HANS-SUB) we note that the INDEX-DIFF factor, which specifically examines tokens in the shared subsequence, captures the setting’s pathological heuristic.

On the constituent OOD set (HANS-CON) we evaluate a factor that examines the attribution on the control words of the premise. For example, for the premise “*Unless the doctors ran, the lawyers encouraged the scientists*” and the hypothesis “*The doctors ran*”, we would consider the attributions on the word “*Unless*”.

**PAWS Factors** We further investigate two intuitive heuristics that are based on the construction of the OOD set. SWAP-AVG uses the average attribution across all swapped tokens and SWAP-MAX-DIFF subtracts the highest magnitude attribution of swapped tokens in the first sentence and the highest magnitude attribution of swapped tokens in the second sentence. For example, for the pair (“*What factors cause a good person to become bad ?*”, “*What factors cause a bad person to become good ?*”), SWAP-AVG would consider the attributions on “*good*” and “*bad*”. SWAP-MAX-DIFF is analogous.

## 6.2 Inoculated Results

We first evaluate models that differ primarily through inoculation, as described in Section 4.1.

Ranking Method	$D_{(O,t)}$		
	PAWS	HANS-SUB	HANS-CON
<b>Baselines</b>			
ACC	88.7	90.6	81.6
CONF	9.2	40.4	52.8
CONF-GT	34.9	20.2	38.9
RANDOM	50.7	51.4	49.6
<b>Factors</b>			
CONST	—	—	<b>87.1</b>
SWAP-MAX-DIFF	76.2	—	—
SWAP-AVG	<b>91.4</b>	—	—
INDEX-DIFF	60.5	<b>91.3</b>	68.6
MAX-DIFF	65.9	69.2	50.6
SUM-DIFF	56.6	60.0	75.2
FIRST-TOK	74.3	50.4	55.6

Table 2: Few-shot heuristic ranking performance on OOD samples for HANS/MNLI and QQP/PAWS, specifically when comparing **inoculated models** (SHAP explanations). We divide rows by baselines, dataset-specific factors, and general factors.

Results are shown using SHAP in Table 2 which we selected through experiments in this setting as being the best performing. The conclusions here differ somewhat from those on MSGS. We note that the ACC baseline remains strong, while CONF is near random. We find that *certain* attribution factors are able to outperform the ACC baseline, with SWAP-AVG the best on PAWS (91.4%), INDEX-DIFF the best on HANS-SUB (91.3%), and CONST on HANS-CON (87.1%).

This shows that even in settings more realistic than MSGS, **the right choice of factor reveals meaningful information about model generalization**. Moreover, the heuristics that work well are those hand-designed for these datasets, confirming that measuring association with a heuristic via a factor may reveal something about performance.

We qualify these results by noting that in a true few-shot setting, there is some uncertainty regarding whether a chosen factor is truly the best one. As a coarse option, we find ACC to be reliable. However, these high-performing factors would still be useful in conjunction with accuracy, or if we had previously validated a factor as ranking models well and we wanted to apply it to rank new models in this domain; the factors will generalize to new models even if they do not generalize to new datasets necessarily.

## 6.3 Architectural Change Results

We further examine our approach when ranking the performance of different pre-trained models

Ranking Method	$D_{(O,t)}$
	HANS/PAWS POOLED
<b>Baseline</b>	
ACCURACY	61.0
GUESS	72.0
<b>Factors</b>	
SET-DEPENDENT	75.5
MAX-DIFF	72.6
INDEX-DIFF	83.0
SUM-DIFF	83.1
FIRST-TOK	69.4

Table 3: Few-shot heuristic ranking performance on OOD samples for HANS/MNLI and QQP/PAWS, specifically when comparing **non-inoculated models** (SHAP explanations), where we take the mean of pairwise accuracies for 3 pairs (for 3 models) on each set.

(RoBERTa, ELECTRA, and DeBERTa).

Table 3 shows that a heuristic GUESS based on the expectation across choosing a *best* model and then randomly guessing consistently with that, gives a strong baseline of 72%. Factors also seem to do well in this setting, with all of the general heuristics outperforming the very low ACC baseline.

This suggests that in few-shot factors are able to capture distributional information that baselines can't. However, to qualify this, given that each set only compares between 3 pairs of models, it's easier for factors to happen upon strong accuracy patterns by chance.

Thus, in Table 4, we analyze this further by showing results on some individual model pairs. (R1, R2 are RoBERTa; E is Electra; D is DeBERTa) R1-E and D-R2, have different architectures, but similar OOD accuracy (see Table 6 in the Appendix). R1-D, E-D and E-R2 are different model types with more distant accuracy. Accuracy values for a single pair on this single dataset therefore only reflect differences across bootstrap samples. What we find in common across these types of pairings is that while some values are close to 50%, including the ACCURACY baseline, each column has several factors achieving very distinct (0%, 100%) accuracy values, consistently differentiating these models. As we note in Figure 4 (Appendix), this pattern of strong distinctions is quite common when different types of models are compared. We further discuss this in Section 7.

## 7 Analysis / Discussion

**Accuracy is reliable, but factors can provide more fine-grained distinctions.** On MSGS,

Ranking Method	Model Pair ( $M_1-M_2$ )				
	R1-E	R1-D	E-R2	D-R2	E-D
<b>Baseline</b>					
ACCURACY	77.4	54.6	55.2	67.2	64.4
<b>Factors</b>					
SWAP-MAX-DIFF	57.6	57.6	70.6	65.6	42.6
SWAP-AVG	93.4	93.4	65.0	47.6	17.0
MAX-DIFF	63.2	63.2	39.2	4.8	0.2
INDEX-DIFF	0	0	99.6	99.6	15.0
SUM-DIFF	99.8	100	0	0	88.0
FIRST-TOK	1.2	100	99.6	0	0

Table 4: SHAP pairwise accuracies, different types of models, PAWS. Model R1 (69.7%), R2 (82.9%), Model E (80.5%), and Model D (71.8%)

where factors beat strong accuracy baselines, we notice that these pairwise accuracies are consistently high. For example, on the MORPH setting, for two models with 95% and 98% accuracy, our factors IRREG is 100% accurate, while the accuracy baseline here is only 58%, as test accuracy on  $D_{(O,t)}$  does not discriminate well between two models with such close overall accuracy.

This holds at the fine-grained pairwise level as well. Figure 5 (also see Figure 6 in appendix) shows the baseline  $D_{(O,t)}$  accuracy against a specific factor's accuracy for each model pair in MSGS. Each datapoint in the scatterplot represents a model pair and a point's vertical distance from the red line represents how much better or worse a given factor does compared to the baseline on a specific pair. We see a regular trend: explanations seem to systematically outperform the baseline across various pairs, with a few significant deviations for low-performing pairs.

These results suggest that explanations can be useful and do add information otherwise missing from accuracy probing alone, especially when the underlying model architecture is held constant. With differing architectures (Figure 6), the problem is made more difficult, and selecting the right factor is less obvious; few-shot accuracy may be more reliable in this setting. Note, however, that these successes from any technique are in spite of us only inspecting 10 examples from the target domain.

**Factors differentiate models strongly, though not always in a way aligned with OOD performance.** Figure 4 and Table 4 both show that factors will often consistently decide in favor of a certain model regardless of the choice of  $D_{(O,t)}$ , especially when dealing with models with different




Figure 4: Distributions of pairwise accuracies on PAWS SHAP non-inoculated, all model pairs (left for accuracy baseline, right for all factors).




Figure 5: LIME pairwise factor against baseline accuracies for MSGs. Additionally see Figure 6.

base architectures. Since ranking accuracy correlates with whether these strong alignments are consistent across a spectrum of models and choose the models with higher OOD performance, the tendency for factors to strongly favor a specific model doesn't necessarily correlate with strong overall performance, but does heavily imply that these factors extract meaningful information about the model from the attributions. Looking close at Table 4, we can see that that even between different model architectures, certain factors are more (INDEX-DIFF) or less (SWAP-MAX-DIFF) capable of making these distinctions.

**Factors as projections of model feature space.** Based on these results, we have evidence that the distributions of attributions are unique to models: in other words, a factor is like a scalar *signature* for a model's feature space with respect to some relevant features. Methods like inoculation, that change a model's behavior in direct ways lead to regular changes in that signature. In these cases,

factors align with OOD performance, which explains why factors are so strong in our inoculated experiments. For our non-inoculated experiments (i.g. ELECTRA vs DeBERTa), the feature spaces are fundamentally different, so factor signatures will still capture these differences, but in a way less aligned with ranking on OOD performance. Future work may be able to expand on these differences and what they tell us beyond OOD performance.

## 8 Related Work

This paper relates to a long line of work on understanding explanations, including explanations' human interpretability (Miller, 2019; Jacovi and Goldberg, 2020; Alqaraawi et al., 2020; Nguyen et al., 2021), explanations' faithfulness and ability to detect shortcuts (Geirhos et al., 2020) or spurious features (Bastings et al., 2021; Madsen et al., 2021; Zhou et al., 2021), and applications to OOD data (Ye and Durrett, 2022; Choi et al., 2022), including papers in the intersection of multiple directions (Adebayo et al., 2022; Kim et al., 2021).

Past work has also investigated performance prediction (Xia et al., 2020; Ye et al., 2021b; Varshney et al., 2022), and using explanations to detect spurious correlations (Kim et al., 2021; Bastings et al., 2021; Adebayo et al., 2022). We are different in that we focus on ranking an array of models which exhibit different levels of generalization abilities, as opposed to giving a binary judgment of whether a model is relying on some shortcuts (Kim et al., 2021; Bastings et al., 2021; Adebayo et al., 2022). In addition, we experiment with tasks having nuanced shortcuts 'in the wild', contrary to synthetically constructed datasets in Bastings et al. (2021). In particular, Adebayo et al. (2022) study the usefulness of explanations in detecting unknown spurious features in an image classification task involving (realistic) possible shortcuts, but find that attributions are ineffective for detecting unknown shortcuts in practice.

## 9 Conclusion

We establish a robust framework for evaluation of fine-grained few-shot prediction of OOD performance, benchmarking approaches in this setting on a range of models. We find that accuracy is a reliable baseline, but intuitive attribution-based factors derived from explanations can sometimes better predict how models will perform in OOD settings, even when they have similar in-domain



performance. We further analyze patterns of our approaches, discovering the potential for factors to represent views of model feature space, leaving further exploration to future work.

## 10 Limitations

There are a large number of explanation techniques and many domains these have been applied to. We focus here on a set of textual reasoning tasks like entailment where spurious correlations have been frequently identified. However, correlations in other settings like medical imaging (Adebayo et al., 2022) could yield different results. We also note that these datasets are all English-language and use English pre-trained models, so different settings may yield different results; additionally, our factors depend on how explanations are normalized between different examples.

Our paper and analysis themselves comment on the limitations of our methodology as well as explanations as a whole: we find that while explanations often can clearly distinguish different models, knowing which factors will do so, or guaranteeing that explanations align with OOD performance, remains difficult.

## Acknowledgments

This work was supported by NSF CAREER Award IIS-2145280, a gift from Salesforce, Inc., and a gift from Adobe. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources used to conduct this research.

## References

- Julius Adebayo, Michael Muelly, Harold Abelson, and Been Kim. 2022. [Post hoc explanations may be ineffective for detecting unknown spurious correlation](#). In *International Conference on Learning Representations*.
- Ahmed Alqaraawi, Martin Schuessler, Philipp Weiß, Enrico Costanza, and Nadia Berthouze. 2020. [Evaluating saliency map explanations for convolutional neural networks: A user study](#). In *Proceedings of the 25th International Conference on Intelligent User Interfaces, IUI '20*, page 275–285, New York, NY, USA. Association for Computing Machinery.
- Jasmijn Bastings, Sebastian Ebert, Polina Zablotskaia, Anders Sandholm, and Katja Filippova. 2021. ["Will You Find These Shortcuts?" A Protocol for Evaluating the Faithfulness of Input Saliency Methods for Text Classification](#). In *arXiv*.
- Jifan Chen and Greg Durrett. 2019. [Understanding Dataset Design Choices for Multi-hop Reasoning](#). 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 4026–4032, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jihye Choi, Jayaram Raghuram, Ryan Feng, Jiefeng Chen, Somesh Jha, and Atul Prakash. 2022. [Concept-based explanations for out-of-distribution detectors](#). In *arXiv*.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training text encoders as discriminators rather than generators](#). In *International Conference on Learning Representations (ICLR)*.
- Matt Gardner, William Merrill, Jesse Dodge, Matthew Peters, Alexis Ross, Sameer Singh, and Noah A. Smith. 2021. [Competency Problems: On Finding and Removing Artifacts in Language Data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1801–1813, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A. Wichmann. 2020. [Shortcut learning in deep neural networks](#). *Nature Machine Intelligence*, 2(11):665–673.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation Artifacts in Natural Language Inference Data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- David Harbecke. 2021. [Explaining natural language processing classifiers with occlusion and language modeling](#). *arXiv preprint arXiv:2101.11889*.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021a. [DeBERTaV3: Improving DeBERTa using ELECTRA-Style Pre-Training with Gradient-Disentangled Embedding Sharing](#). In *arXiv ePrint 2111.09543*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021b. [Deberta: Decoding-enhanced bert with disentangled attention](#). In *International Conference on Learning Representations*.
- Dieuwke Hupkes, Mario Giulianelli, Verna Dankers, Mikel Artetxe, Yanai Elazar, Tiago Pimentel, Christos Christodoulopoulos, Karim Lasri, Naomi Saphra, Arabella Sinclair, Dennis Ulmer, Florian Schottmann, Khuyagbaatar Batsuren, Kaiser Sun, Koustuv Sinha,

- Leila Khalatbari, Maria Ryskina, Rita Frieske, Ryan Cotterell, and Zhijing Jin. 2022. [State-of-the-art generalisation research in nlp: a taxonomy and review](#).
- Shankar Iyer, Nikhil Dandekar, and Kornel Csernai. 2017. First Quora dataset release: Question pairs.
- 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.
- Robin Jia and Percy Liang. 2017. [Adversarial Examples for Evaluating Reading Comprehension Systems](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2021–2031, Copenhagen, Denmark. Association for Computational Linguistics.
- Brendan Kennedy, Xisen Jin, Aida Mostafazadeh Davani, Morteza Dehghani, and Xiang Ren. 2020. [Contextualizing Hate Speech Classifiers with Post-hoc Explanation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5435–5442, Online. Association for Computational Linguistics.
- Joon Sik Kim, Gregory Plumb, and Ameet Talwalkar. 2021. [Sanity simulations for saliency methods](#). *CoRR*, abs/2105.06506.
- Frederick Liu and Besim Avci. 2019. [Incorporating Priors with Feature Attribution on Text Classification](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6274–6283, Florence, Italy. Association for Computational Linguistics.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019a. [Inoculation by Fine-Tuning: A Method for Analyzing Challenge Datasets](#). 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 2171–2179, Minneapolis, Minnesota. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [Roberta: A robustly optimized bert pretraining approach](#). In *arXiv*.
- Scott M. Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NeurIPS’17*, page 4768–4777, Red Hook, NY, USA. Curran Associates Inc.
- Andreas Madsen, Nicholas Meade, Vaibhav Adlakha, and Siva Reddy. 2021. Evaluating the Faithfulness of Importance Measures in NLP by Recursively Masking Allegedly Important Tokens and Retraining. In *arXiv ePrint 2110.08412*.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Tim Miller. 2019. [Explanation in artificial intelligence: Insights from the social sciences](#). *Artificial Intelligence*, 267:1–38.
- Giang Nguyen, Daeyoung Kim, and Anh Nguyen. 2021. [The effectiveness of feature attribution methods and its correlation with automatic evaluation scores](#). In *Advances in Neural Information Processing Systems*.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis Only Baselines in Natural Language Inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Danish Pruthi, Rachit Bansal, Bhuwan Dhingra, Livio Baldini Soares, Michael Collins, Zachary C. Lipton, Graham Neubig, and William W. Cohen. 2022. [Evaluating Explanations: How Much Do Explanations from the Teacher Aid Students?](#) *Transactions of the Association for Computational Linguistics*, 10:359–375.
- 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*.
- 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 Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 2662–2670.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. 2017. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML’17*, page 3319–3328. JMLR.org.
- Liyan Tang, Dhruv Rajan, Suyash Mohan, Abhijeet Pradhan, R. Nick Bryan, and Greg Durrett. 2021. [Making Document-Level Information Extraction Right for the Right Reasons](#). *arXiv*.
- Neeraj Varshney, Swaroop Mishra, and Chitta Baral. 2022. Ildae: Instance-level difficulty analysis of evaluation data. *ArXiv*, abs/2203.03073.
- Alex Warstadt, Yian Zhang, Xiaocheng Li, Haokun Liu, and Samuel R. Bowman. 2020. [Learning which features matter: RoBERTa acquires a preference for linguistic generalizations \(eventually\)](#). In *Proceedings of the 2020 Conference on Empirical Methods*

in *Natural Language Processing (EMNLP)*, pages 217–235, Online. Association for Computational Linguistics.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.

Mengzhou Xia, Antonios Anastasopoulos, Ruochen Xu, Yiming Yang, and Graham Neubig. 2020. [Predicting performance for natural language processing tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8625–8646, Online. Association for Computational Linguistics.

Huihan Yao, Ying Chen, Qinyuan Ye, Xisen Jin, and Xiang Ren. 2021. [Refining Language Models with Compositional Explanations](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 8954–8967. Curran Associates, Inc.

Xi Ye and Greg Durrett. 2022. [Can Explanations Be Useful for Calibrating Black Box Models?](#) In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6199–6212, Dublin, Ireland. Association for Computational Linguistics.

Xi Ye, Rohan Nair, and Greg Durrett. 2021a. [Connecting Attributions and QA Model Behavior on Realistic Counterfactuals](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5496–5512, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Zihuiwen Ye, Pengfei Liu, Jinlan Fu, and Graham Neubig. 2021b. [Towards more fine-grained and reliable NLP performance prediction](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3703–3714, Online. Association for Computational Linguistics.

Yuan Zhang, Jason Baldridge, and Luheng He. 2019. [PAWS: Paraphrase adversaries from word scrambling](#). 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 1298–1308, Minneapolis, Minnesota. Association for Computational Linguistics.

Yilun Zhou, Serena Booth, Marco Tulio Ribeiro, and Julie Shah. 2021. [Do feature attribution methods correctly attribute features?](#) In *eXplainable AI approaches for debugging and diagnosis*.

	VERB	MORPH	ADJECT
NO-INOC	12.0	95.0	51.0
2L	99.0	98.0	99.2
2S	0.0	0.0	0.0
2L 1S	80.0	68.0	73.0
1L 2S	33.0	57.0	32.0
2L 2S	53.7	49.0	56.0

Table 5: MSGS accuracies of various inoculated models.

## A Details of Inoculation

One of the methods we used to obtain models with different performances on the OOD sets was inoculation (Liu et al., 2019a), which involves fine-tuning or further fine-tuning models on small amounts or batches of OOD data alongside in-domain data to bring model performance on OOD sets up.

**MSGS** We borrow notation from Warstadt et al. (2020). Most of the fine-tuning data is ambiguous data that doesn’t test the spurious correlation, but we add in small percentages of non-ambiguous data where the label favors either the surface or linguistic generalization, tilting the model in that direction. Here, for each set (VERB, MORPH, ADJECT), we used the following inoculation splits. Linguistic (L) and surface(S) are the features that the inoculation data would favor: 2% L, 2% S, mixed (2% L, 1% S), (1% L 2% S), (2% L, 2% S), in addition to no inoculation. The results on  $D_O$  are present in Table 5.

**HANS** Specific inoculation results for RoBERTa-large are present in Table 6. We additionally use MNLI pre-trained ELECTRA and DeBERTa models from huggingface. These performance details are also located in Table 6.

**PAWS** We used several inoculation techniques to get a variable number of models here. For our RoBERTa-base model, we start with the base model (35% OOD accuracy) and fine-tune it with  $D_T$  data with 2% of the data having  $D_O$  data mixed in. We trained this over several epochs to get models with 82.8% and 90.8% accuracy on  $D_O$ . We also tried fine-tuning our 35% model on batches of pure  $D_O$  data to get a model with 69% accuracy. For our ELECTRA and DeBERTa models, we use similar batch-only inoculation (fine-tuning on batches of only OOD data). More details are present in Table 6.

Tables 7-8 contain the same information as Table 2, but for the other 2 studied explanation tech-

nique.

## B Bootstrapping Details

We now describe our process for bootstrapping and evaluating the capability of explanations in our setting.

For a sampled population of examples from the  $D_O$  set, for the  $m$  models that we’re examining at a time, we generate explanations for each of the  $m$  models on all of the sampled population. We then repeatedly take a sample with replacement (500 times) of 10 examples  $D_{O,t}$  each, where we have  $500 \times 10 \times m$  total explanations we want to examine. We calculate factors for each of the 10 explanations in each  $D_{O,t}$  sample and pool them to get a list of factor metrics for the  $D_{O,t}$ , one for each model.

For each pair, we then look at the ground-truth  $D_O$  ranking for models and their respective factor metrics, getting *successes* where these match, and *failures* otherwise. When we average these accuracies across our 500 bootstrap samples, we get pairwise distributions (the distribution of successes vs failures on a sample for a given pair), which we can further aggregate to get few-shot accuracies.

Note, in practice, to prevent variance from run-to-run, we fix the population of 500  $D_{O,t}$ s, but we validated that re-running on new sampled populations didn’t impact any numbers greatly. Though we tried using several (5, 10, 20)  $D_{O,t}$  sizes, we decided to use the probe size of 10 as a realistic probe size for our setting, which wouldn’t be burdensome to hand-craft in practice.

Our methodology can be run quickly in a post-hoc manner as many times as needed on top of a population of the necessary explanations.

## C Additional Plots

Figure 4 shows additional information about the distribution of pairwise accuracies between different model architectures.

## D Reproducibility

### D.1 Computing Infrastructure

All experiments were conducted on a desktop with 2 NVIDIA 1080 Ti (11 GB) and 1 NVIDIA Titan Xp (12 GB).

### D.2 Runtimes

For PAWS and MSGS fine-tuned models, we fine-tuned for roughly 1 GPU hour per model. Since




Figure 6: SHAP pairwise factor compared to ACC for HANS-CON and PAWS. Each point represents a factor accuracy (y-axis) for a pair of models in comparison to ACC (x-axis) for the same pair. Points above the red  $y = x$  line represent factors outperforming the accuracy baseline. CONST and DIFF-SUM are for HANS-CON, SWAP-AVG and SWAP-MAX-DIFF are for PAWS

HANS models were trained for very few steps, their training time is inconsequential. Generating attributions required for numerical evaluation took less than 6 GPU hours.

### D.3 Dataset Details

We used datasets in the JSONL format. We simplified all our dataset settings to binary classification for simplicity, and used data directly from the downloads made available in the original papers.

Dataset	OOD Performance	Huggingface Model Name	LR	Warmup	Steps
HANS	99.8/99.4	roberta-large-mnli	$1e^{-5}$	500	150
HANS	96.7/97.6	roberta-large-mnli	$1e^{-5}$	500	100
HANS	87.1/70.1	roberta-large-mnli	$1e^{-5}$	500	75
HANS	79.5/62.5	roberta-large-mnli	$1e^{-5}$	500	50
HANS	69.9/58.7	roberta-large-mnli	$1e^{-5}$	500	25
HANS	66.8/57.8	roberta-large-mnli	—	—	—
HANS	63.5/72.5	howey/electra-base-mnli	—	—	—
HANS	62.7/65.7	MoritzLaurer/DeBERTa-v3-base-mnli	—	—	—
MSGS	Table 5	roberta-base	$1e^{-5}$	600	6000
PAWS	90.8	roberta-base	$1e^{-5}$	1200	12000
PAWS	82.8	roberta-base	$1e^{-5}$	1200	12000
PAWS	69.0	roberta-base	$1e^{-5}$	1200	7600
PAWS	35.0	roberta-base	$1e^{-5}$	1200	12000
PAWS	80.5	google/electra-base-discriminator	$1e^{-5}$	1200	7700
PAWS	71.8	microsoft/deberta-base	$1e^{-5}$	1200	7600

Table 6: Architecture details for our experiments. “Steps” indicates the number of gradient updates from the specified dataset that are applied to the model. For HANS models, performance is on HANS-SUB/HANS-CON. For all models, small batch sizes were used, with weight decay of 0.1.

	PAWS	HANS-SUB	HANS-CON
<b>Baselines</b>			
ACCURACY	88.7	90.6	81.6
CONFIDENCE	9.2	40.4	52.8
RANDOM	50.7	51.4	49.6
<b>Explanations</b>			
CONST	—	—	79.4
SWAP-MAX-DIFF	80.6	—	—
SWAP-AVG	98.2	—	—
MAX-DIFF	70.1	67.2	58.3
INDEX-DIFF	70.5	88.5	67.2
SUM-DIFF	53.9	59.1	60.4
FIRST-TOK	55.4	51.0	81.3

Table 7: LIME version of Table 3

	PAWS	HANS-SUB	HANS-CON
<b>Baselines</b>			
ACCURACY	88.7	90.6	81.6
CONFIDENCE	9.2	40.4	52.8
RANDOM	50.7	51.4	49.6
<b>Explanations</b>			
CONST	—	—	79.2
SWAP-MAX-DIFF	84.3	—	—
SWAP-AVG	85.6	—	—
MAX-DIFF	86.9	55.2	69.9
INDEX-DIFF	51.4	85.8	53.4
SUM-DIFF	51.6	77.4	68.1
FIRST-TOK	64.0	69.7	59.2

Table 8: Tokig numbers for Table 3

ID Set	OOD Set	$D_O$ Size	$D_{(O,t)}$ Size
MSGS	MORPH	10000	10
	VERB	10000	10
	ADJECT	10000	10
MNLI	HANS-SUB	10000	10
	HANS-CON	10000	10
QQP	PAWS	677	10

Table 9: Information regarding our considered datasets. For all datasets, the bootstrap sample size is fixed at 10.