

Is My Model Using the Right Evidence? Systematic Probes for Examining Evidence-Based Tabular Reasoning

Vivek Gupta¹, Riyaz A. Bhat², Atreya Ghosal³,
Manish Shrivastava³, Maneesh Singh², Vivek Srikumar¹

¹University of Utah, USA, ²Verisk Inc., India, ³IIT-Hyderabad, India
{vgupta, svivek}@cs.utah.edu, {riyaz.bhat, maneesh.singh}@verisk.com,
{atreyyee.ghosal@research, m.shrivastava}@iiit.ac.in

Abstract

Neural models command state-of-the-art performance across NLP tasks, including ones involving “reasoning”. Models claiming to reason about the evidence presented to them should attend to the correct parts of the input while avoiding spurious patterns therein, be self-consistent in their predictions across inputs, and be immune to biases derived from their pre-training in a nuanced, context-sensitive fashion. *Do the prevalent *BERT-family of models do so?* In this paper, we study this question using the problem of reasoning on tabular data. Tabular inputs are especially well-suited for the study—they admit systematic probes targeting the properties listed above. Our experiments demonstrate that a RoBERTa-based model, representative of the current state-of-the-art, fails at reasoning on the following counts: it (a) ignores relevant parts of the evidence, (b) is over-sensitive to annotation artifacts, and (c) relies on the knowledge encoded in the pre-trained language model rather than the evidence presented in its tabular inputs. Finally, through inoculation experiments, we show that fine-tuning the model on perturbed data does not help it overcome the above challenges.

1 Introduction

The problem of understanding tabular or semi-structured data is a challenge for modern NLP. Recently, Chen et al. (2020b) and Gupta et al. (2020) have framed this problem as a natural language inference question (NLI, Dagan et al., 2013; Bowman et al., 2015, *inter alia*) via the TabFact and the INFO TABS datasets, respectively. The tabular version of the NLI task seeks to determine whether a *tabular* premise entails or contradicts a textual hypothesis, or is unrelated to it.

One strategy for such tabular reasoning tasks relies on the successes of contextualized representations (e.g., Devlin et al., 2019; Liu et al., 2019b) for the sentential version of the problem. Tables are flattened into artificial sentences using heuristics to be processed by these models. Surprisingly, even this naïve strategy leads to high predictive accuracy, as shown not only by the introductory papers but also by related lines of recent work (e.g., Eisenschlos et al., 2020; Yin et al., 2020).

In this paper, we ask: *Do these seemingly accurate models for tabular inference effectively use and reason about their semi-structured inputs?* While “reasoning” can take varied forms, a model that claims to do so should at least ground its outputs on the evidence provided in its inputs. Concretely, we argue that such a model should (a) be self-consistent in its predictions across controlled variants of the input, (b) use the evidence presented to it, and the right parts thereof, and, (c) avoid being biased *against* the given evidence by knowledge encoded in the pre-trained embeddings.

Corresponding to these three properties, we identify three dimensions on which to evaluate a tabular NLI system: robustness to annotation artifacts, relevant evidence selection, and robustness to counterfactual changes. We design systematic probes that exploit the semi-structured nature of the premises. This allows us to semi-automatically construct the probes and to unambiguously define the corresponding expected model response. These probes either introduce controlled edits to the premise or the hypothesis, or to both, thereby also creating counterfactual examples. Experiments reveal that despite seemingly high test set accuracy, a model based on RoBERTa (Liu et al., 2019b), a good representative of BERT derivative models, is far from being reliable. Not only

does it ignore relevant evidence from its inputs, it also relies excessively on annotation artifacts, in particular the sentence structure of the hypothesis, and pre-trained knowledge in the embeddings. Finally, we found that attempts to inoculate the model (Liu et al., 2019a) along these dimensions degrades its overall performance.

The rest of the paper is structured as follows. §2 introduces the Tabular NLI task, while §3 articulates the need for probing evidence-based tabular reasoning in extant high-performing models. §4–§6 detail the probes designed for such an examination and the results thereof while §7 analyzes the impact of inoculation to aforementioned challenges through model fine-tuning. §8 presents the main takeaways and contextualization in the related art. §9 provides concluding remarks and indicates future directions of work.¹

2 Preliminaries: Tabular NLI

Tabular natural language inference is a task similar to standard NLI in that it examines if a natural language hypothesis can be derived from the given premise. Unlike standard NLI, where the evidence is presented in the form of sentences, the premises in tabular NLI are semi-structured tables that may contain both text and data.

Dataset Recently, datasets such as TabFact (Chen et al., 2020b) and INFO TABS (Gupta et al., 2020), and also shared tasks such as SemEval 2021 Task 9 (Wang et al., 2021a) and FEVEROUS (Aly et al., 2021), have sparked interest in tabular NLI research. In this study, we use the INFO TABS dataset for our investigations.

INFO TABS consists of 23,738 premise-hypothesis pairs, whose premises are based on Wikipedia infoboxes. Unlike TabFact, which only contains **ENTAIL** and **CONTRADICT** hypotheses, INFO TABS also includes **NEUTRAL** ones. Figure 1 shows an example table from the dataset with four hypotheses, which will be our running example.

The dataset contains 2,540 distinct infoboxes representing a variety of domains. All hypotheses were written and labeled by Amazon MTurk workers. The tables contain a *title* and two columns, as shown in the example. Since each row takes the form of a key-value pair, we will refer to the

¹The dataset and the scripts used for our analysis are available at <https://tabprobe.github.io>.

Breakfast in America	
Released	29 March 1979
Recorded	May–December 1978
Studio	The Village Recorder (Studio B) in Los Angeles
Genre	pop ; art rock ; soft rock
Length	46:06
Label	A&M
Producer	Peter Henderson, Supertramp

- H1: Breakfast in America is a pop album with a length of 46 minutes.
H2: Breakfast in America was released at the end of 1979.
H3: Most of Breakfast in America was recorded in the last month of 1978.
H4: Breakfast in America has 6 tracks.

Figure 1: An example of a tabular premise from INFO TABS. The hypotheses H1 is entailed by it, H2 contradicts it, and H3, H4 are neutral (i.e., neither entailed nor contradictory).

elements in the left column as the *keys*, and the right column provides the corresponding *values*.

In addition to the usual train and development sets, INFO TABS includes three test sets, α_1 , α_2 , and α_3 . The α_1 set represents a standard test set that is both topically and lexically similar to the training data. In the α_2 set, hypotheses are designed to be lexically adversarial, and the α_3 tables are drawn from topics not present in the training set. We will use all three test sets for our analysis.

Models over Tabular Premises Unlike standard NLI, which can use off-the-shelf pre-trained contextualized embeddings, the semi-structured nature of premises in tabular NLI necessitates a different modeling approach.

Following Chen et al. (2020b), tabular premises are flattened into token sequences that fit the input interface of such models. While different flattening strategies exist in the literature, we adopt the *Table as a Paragraph* strategy of Gupta et al. (2020), where each row is converted to a sentence of the form “The key of `title` is `value`”. This seemingly naïve strategy, with RoBERTa-large embeddings (RoBERTa_L henceforth), achieved the highest accuracy in the original work, shown in

Model	dev	α_1	α_2	α_3
Human	79.78	84.04	83.88	79.33
Hypothesis Only	60.51	60.48	48.26	48.89
RoBERTa _L	75.55	74.88	65.55	64.94
5xCV	73.59 _(2,3)	72.41 _(1,4)	63.02 _(1,9)	61.82 _(1,4)

Table 1: Results of the *Table as a Paragraph* strategy on INFOABS subsets with RoBERTa_L model, hypothesis-only baseline and majority human agreement. The first three rows are reproduced from Gupta et al. (2020). The last row represents the average performances (and standard deviations as subscripts) using models obtained via five-fold cross validation (5xCV).

Table 1.² The table also shows the hypothesis-only baseline (Poliak et al., 2018; Gururangan et al., 2018) and human agreement on the labels.³

To study the stability of the models to variations in the training data, we performed 5-fold cross validation (5xCV). An average cross validation accuracy of 73.53% with a standard deviation of 2.73% was observed on the training set which is close to the performance on the α_1 test set (74.88%). In addition, we also evaluated performance on the development and test sets. The penultimate row of Table 1 presents the performance for the model trained on the entire training data, and the last row presents the performance of the 5xCV models. The results demonstrate that model performance is reasonably stable to variations in the training set.

3 Reasoning: An Illusion?

Given the surprisingly high accuracies in Table 1, especially on the α_1 test dataset, can we conclude that the RoBERTa-based model reasons effectively about the evidence in the tabular input to make its inference? That is, does it arrive at its answer via a sound logical process that takes into account all available evidence along with com-

²Other flattening strategies have similar performance (Gupta et al., 2020).

³Preliminary experiments on the development set showed that RoBERTa_L outperformed other pre-trained embeddings. We found that BERT_B, RoBERTa_B, BERT_L, ALBERT_B, and ALBERT_L reached development set accuracies of 63.0%, 67.23%, 69.34%, 70.44%, and 70.88%, respectively. While we have not replicated our experiments on these other models due to a prohibitively high computational cost, we expect the conclusions to carry over to these other models as well.

mon sense knowledge? Merely achieving high accuracy is not sufficient evidence of reasoning: The model may arrive at the right answer for the wrong reasons leading to improper and inadequate generalization over unseen data. This observation is in line with the recent work pointing out that the high-capacity models we use may be relying on spurious correlations (e.g., Poliak et al., 2018).

“Reasoning” is a multi-faceted phenomenon, and fully characterizing it is beyond the scope of this work. However, we can probe for the *absence* of evidence-grounded reasoning via model responses to carefully constructed inputs and their variants. The guiding premise for this work is:

Any “evidence-based reasoning” system should demonstrate expected, predictable behavior in response to controlled changes to its inputs.

In other words, “reasoning failures” can be identified by checking if a model deviates from expected behavior in response to controlled changes to inputs. We note that this strategy has been either explicitly or implicitly employed in several lines of recent work (Ribeiro et al., 2020; Gardner et al., 2020). In this work, we instantiate the above strategy along three specific dimensions, briefly introduced here using the running example in Figure 1. Each dimension is used to define several concrete probes that subsequent sections detail.

1. Avoiding Annotation Artifacts A model should not rely on spurious lexical correlations. In general, it should not be able to infer the label using only the hypothesis. Lexical differences in closely related hypotheses should produce predictable changes in the inferred label. For example, in the hypothesis H2 of Figure 1, if the token “end” is replaced with “start”, the model prediction should change from **CONTRADICT** to **ENTAIL**.

2. Evidence Selection A model should use the correct evidence in the premise for determining the hypothesis label. For example, ascertaining that the hypothesis H1 is entailed requires the *Genre* and *Length* rows of Figure 1. When a relevant row is removed from a table, a model that predicts the **ENTAIL** or the **CONTRADICT** label should predict the **NEUTRAL** label. When an irrelevant row is removed, it should not change its prediction from **ENTAIL** to **NEUTRAL** or vice versa.

3. Robustness to Counterfactual Changes A model’s prediction should be *grounded* in the provided information even if it contradicts the real world, that is to counterfactual information. For example, if the month of the *Released* date changed to “December”, then the model should change the label of H2 in Figure 1 to **ENTAIL** from **CONTRADICT**. Since this information about release date contradicts the real world, the model cannot rely on its pre-trained knowledge, say, from Wikipedia. For the model to predict the label correctly, it needs to reason with the information in the table as the primary evidence. Although the importance of pre-trained knowledge cannot be overlooked, it must not be at the expense of primary evidence.

Further, there are certain pieces of information in the premise (irrelevant to the hypothesis) that do not impact the outcome, making the outcome *invariant* to these changes. For example, deleting irrelevant rows from the premise should not change the model’s predicted label. Contrary to this is the relevant information (“evidence”) in the premise. Changing these pieces of information should vary the outcome in a predictable manner, making the model *covariant* with these changes. For example, deleting relevant evidence rows should change the model’s predicted label to **NEUTRAL**.

The three dimensions above are not limited to tabular inference. They can be extended to other NLP tasks, such as reading comprehension as well as the standard sentential NLI. However, directly checking for such properties would require considerable labeled data—a big practical impediment. Fortunately, in the case of tabular inference, the (in-/co-)variants associated with these dimensions allow controlled and semi-automatic edits to the inputs leading to predictable variation of the expected output. This insight underlies the design of probes using which we examine the robustness of the reasoning employed by a model performing tabular inference. As we will see in the following sections, highly effective and precise probes can be designed without extensive annotation.

4 Probing Annotation Artifacts

Can a model make inference about a hypothesis without a premise? It is natural to answer in the negative in general. (Of course, certain hypotheses may admit strong priors, e.g., tautologies.)

Preliminary experiments by Gupta et al. (2020) on INFO TABS, however, reveal that a model trained just on hypotheses performs surprisingly well on the test data. This phenomenon, an inductive bias entirely predicated on the hypotheses, is called *hypothesis bias*. Models for other NLI tasks have been similarly shown to exhibit hypothesis bias, whereby the models learn to rely on spurious correlations between patterns in the hypotheses and corresponding labels (Poliak et al., 2018; Gururangan et al., 2018; Geva et al., 2019, and others). For example, negations are observed to be highly correlated with contradictions (Niven and Kao, 2019).

To better characterize a model’s reliance on such artifacts, we perform controlled edits to hypotheses without altering associated premises. Unlike the α_2 set, which includes minor changes to function words, we aim to create more sophisticated changes by altering content expressions or noun phrases in a hypothesis. Two possible scenarios arise where a hypothesis alteration, without a change in the premise, either (a) leads to a change in the label (i.e., the label covaries with the variation in the hypothesis), or (b) does not induce a label change (i.e., the label is invariant to the variation in the hypothesis).

In INFO TABS, a set of reasoning categories are identified to characterize the relationship between a tabular premise and a hypothesis. We use a subset of these, listed below, to perform controlled changes in the hypotheses.

- **Named Entities:** such as *Person, Location, Organization*;
- **Nominal modifiers:** nominal phrases or clauses;
- **Negation:** markers such as *no, not*;
- **Numerical Values:** numeric expressions representing *weights, percentages, areas*;
- **Temporal Values:** Date and Time; and,
- **Quantifiers:** like *most, many, every*.

Although we can easily track these expressions in a hypothesis using tools like entity recognizers and parsers, it is non-trivial to automatically modify them with a predictable change on the hypothesis label. For example, some label changes can only be controlled if the target expression in

Preposition	Upward Monotonicity	Downward Monotonicity
over	CONTRADICT	ENTAIL
under	ENTAIL	CONTRADICT
more than	CONTRADICT	ENTAIL
less than	ENTAIL	CONTRADICT
before	ENTAIL	CONTRADICT
after	CONTRADICT	ENTAIL

Table 2: Monotonicity properties of prepositions.

the hypothesis is correctly aligned with the facts in the premise. Such cases include **CONTRADICT** to **ENTAIL**, and **NEUTRAL** to **CONTRADICT** or **ENTAIL**, which are difficult without extensive expression-level annotations. Nonetheless, in several cases, label changes can be deterministically known even with imprecise changes in the hypothesis. For example, we can convert a hypothesis from **ENTAIL** to **CONTRADICT** by replacing a named entity in the hypothesis with a random entity of the same type.

Hence we follow the following strategy: (a) We avoid perturbations involving the **NEUTRAL** label altogether, as they often need changes in the premise (table) as well. (b) We generate all label-preserving and some label-flipping transformations automatically using the approach described below. (c) We annotate the **CONTRADICT** to **ENTAIL** label-flipping perturbations manually.

Automatic Generation of Label-preserving Transformations To automatically perturb hypotheses, we leverage the syntactic structure of a hypothesis and the monotonicity properties of function words like prepositions. First, we perform syntactic analysis of a hypothesis to identify named entities and their relations to title expressions via dependency paths.⁴ Then, based on the entity type, we either substitute or modify them. Named entities such as person names and locations are substituted with entities of the same type. Expressions containing numbers are modified using the monotonicity property of the prepositions (or other function words) governing them in their corresponding syntactic trees.

Given the monotonicity property of a preposition (see Table 2), we modify its governing numerical expression in a hypothesis in the same order to preserve the hypothesis label. Consider hypothesis H5 in Figure 2, which contains a preposition (*over*) with upward monotonicity. Because of upward monotonicity, we can increase the number of hours in H5 without altering the label.

⁴We used spaCy v2.3.2 for the syntactic analysis.

Bridesmaids	
Length	125 minutes

H5: Bridesmaids has a running time of over 3 hrs.

Figure 2: Hypothesis H5 contradicts the premise.

Type of Modification	Perturbed Hypothesis
Nominal Modifier	H1 ^E _E : Breakfast in America <i>which was produced by Pert Handerson</i> is a pop album of 46 minutes length.
Temporal Expression Negation	H1 ^E _C : Breakfast in America is a pop album with a length of <i>56</i> minutes. H2 ^C _E : Breakfast in America was <i>not</i> released towards the end of 1979.
Temporal Expression	H2 ^C _C : Breakfast in America was released towards the end of <i>1989</i> .

Table 3: Example hypothesis perturbations for the running example from Figure 1. The *red italicized* text represents changes. Superscripts *E/C* represent gold **ENTAIL** and **CONTRADICT** labels, and subscripts *E/C* represent new labels.

Manual Annotation of Label-flipping Transformations Note that in the above example, modifying the numerical expression in the reverse direction (e.g., decreasing the number of hours) does not guarantee a label flip. We need to know the premise to be accurate. During the experiments, we observed that a large step (half/twice the actual number) suffices in most cases. We used this heuristic and manually curated the erroneous cases. Additionally, all the cases of **CONTRADICT** to **ENTAIL** label-flipping perturbations were annotated manually.⁵

We generated 2,891 perturbed examples from the α_1 set with 1,203 instances preserving the label and 1,688 instances flipping it. We also generated 11,550 examples from the *Train* set, with 4,275 preserving and 7,275 flipping the label. Some example perturbations using different types of expressions are listed in Table 3. It should be noted that there may not be a one-to-one correspondence between the gold and perturbed examples, as a hypothesis may be perturbed numerous times or not at all. As a result, in order for the results to be comparable, a single perturbed example must be sampled for each gold example: we sampled 967 from the α_1 set and 4,274 from the *Train* set.

⁵Annotation done by an expert well versed in the NLI task.

Model	Original <i>mean</i> (<i>stdev</i>)	Label Preserved	Label Flipped
<i>Train Set (w/o NEUTRAL)</i>			
Prem+Hypo	99.44 _(0.06)	92.98 _(0.20)	53.92 _(0.28)
Hypo-Only	96.39 _(0.13)	70.23 _(0.35)	19.23 _(0.27)
α_1 Set (w/o NEUTRAL)			
Prem+Hypo	68.94 _(0.76)	69.56 _(0.77)	51.48 _(0.86)
Hypo-Only	63.52 _(0.75)	60.27 _(0.85)	31.02 _(0.63)

Table 4: Results of the Hypothesis-only model and Prem+Hypo model on the gold and perturbed hypotheses.

Results and Analysis We tested the hypothesis-only and full models (both trained on the original *Train* set) on the perturbed examples, without subsequent fine-tuning on the perturbed examples.⁶ The results are presented in Table 4, with each cell representing the average accuracy and standard deviation (subscript) across 100 samplings, with 80% of the data selected at random in each sampling.

We note that the performance degrades substantially in both label-preserved and flipped settings when the model is trained on just the hypotheses. When labels are flipped after perturbations, the decrease in performance (averaged across both models) is about 25% and 61% points, on the α_1 set and *Train* set, respectively. However, for the full model, perturbations that retain the hypothesis label have little effect on model performance.

The contrast in the performance drop between the label-preserved and label-flipped cases suggests that changes to the content expressions have little effect on the model’s original predictions. Interestingly, the predictions are invariant to changes to function words as well, as per results on α_2 in Gupta et al. (2020). This suggests that the model might be more prone to changes to the template or structure of a hypothesis than its lexical makeup. Consequently, a model that relies on correlations between the hypothesis structure and the label is expected to suffer on the label-flipped cases. In case of label-preserving perturbations of similar kind, structural correlations between the hypothesis and the label are retained leading to minimal drop in model performance.

The results of the hypothesis-only model on the *Train* set may appear slightly surprising at first. However, given that the model was trained

⁶We analyze the impact of fine-tuning on perturbed examples in §7.

on this dataset, it seems reasonable to assume that the model has “overfit” to the training data. Therefore, the model is expected to be vulnerable even to slight label-preserving modifications to the examples it was trained on, leading to the huge drop of 26%. In the same setting, for the α_1 set the performance drop is lesser, namely, about 3%.

Taken together, we can conclude from these results that the model ignores the information in the hypotheses (thereby perhaps also the aligned facts in the premise), and instead relies on irrelevant structural patterns in the hypotheses.

5 Probing Evidence Selection

Predictions of an NLI model should primarily be based on the evidence in the premise, that is, on the facts relevant to the hypothesis. For a tabular premise, rows containing the evidence necessary to infer the associated hypothesis are called relevant rows. Short-circuiting the evidence in relevant rows for inference using annotation artifacts as suggested in §4 or other spurious artifacts in irrelevant rows of the table is expected to lead to poor generalization over unseen data.

To better understand the model’s ability to select evidence in the premise, we use two kinds of controlled edits: (a) automatic edits without any information about relevant rows, and (b) semi-automatic edits using knowledge of relevant rows via manual annotation. The rest of the section goes over both scenarios in detail. All experiments in this section use the full model that is trained on both premises and their associated hypotheses.

5.1 Automatic Probing

We define four kinds of table modifications that are agnostic to the relevance of rows to a hypothesis: (a) *row deletion*, (b) *row insertion*, (c) *row-value update*, that is, changing existing information, and (d) *row permutation*, that is, re-ordering rows. Each modification allows certain desired (valid) changes to model predictions.⁷ We examine below the case of row deletion in detail and refer the reader to the Appendix for the others.

Row deletion should lead to the following desired effects: (a) If the deleted row is relevant to the hypothesis (e.g., *Length* for H1), the model

⁷In performing these modifications, we made sure that the modified table does not become inconsistent or self-contradicting.

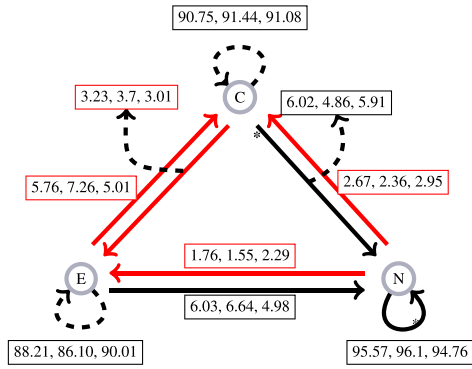


Figure 3: Changes in model predictions after automatic row deletion. Directed edges are labeled with transition percentages from the source node label to the target node label. The number triple corresponds to α_1 , α_2 , and α_3 test sets, respectively and, for each source node, adds up to 100% over the outgoing edges. Red lines represent invalid transitions. Dashed and solid black lines represent valid transitions for irrelevant and relevant row deletion respectively. * represents valid transitions with either row deletions.

prediction should change to **NEUTRAL**. (b) If the deleted row is irrelevant (e.g., *Producer* for H1), the model should retain its original prediction. **NEUTRAL** predictions should remain unaffected by row deletion.

Results and Analysis We studied the impact of *row deletion* on the α_1 , α_2 , and α_3 test sets. Figure 3 shows aggregate changes to labels after row deletions as a directed labeled graph. The nodes in this graph represent the three labels in INFO TABS, and the edges denote transitions after *row deletion*. The source and end nodes of an edge represent predictions before and after the modification.

We see that the model makes invalid transitions in all three datasets. Table 5 summarizes the invalid transitions by aggregating them over the label originally predicted by the model. The percentage of invalid transitions is higher for **ENTAIL** predictions than for **CONTRADICT** and **NEUTRAL**. After row deletion, many **ENTAIL** examples are incorrectly transitioning to **CONTRADICT** rather than to **NEUTRAL**. The opposite trend is observed for the **CONTRADICT** predictions.

As with row deletion, the model exhibits invalid responses to other row modifications listed in the beginning of this section, like row insertion. Surprisingly, the performance degrades due to row permutations as well, suggesting some form of position bias in the model. On the positive

Dataset	α_1	α_2	α_3	Average
ENTAIL	5.76	7.26	5.01	6.01
NEUTRAL	4.43	3.91	5.24	4.53
CONTRADICT	3.23	3.70	3.01	3.31
Average	4.47	4.96	4.42	–

Table 5: Percentage of invalid transitions after row deletion. For an ideal model, all these numbers should be zero.

side, the model mostly retains its predictions on row-value update operations. We refer the reader to the Appendix for more details.

5.2 Manual Probing

Row modification for automatic probing in §5.1 is agnostic to the relevance of the row to a given hypothesis. Since only a few rows (one or two) are relevant to the hypothesis, the probing skew towards hypothesis-unrelated rows weakens the investigations into the evidence-grounding capability of the model. Knowing the relevance of rows allows for the creation of stronger probes. For example, if a relevant row is deleted, the **ENTAIL** and **CONTRADICT** predictions should change to **NEUTRAL**. (Recall that after deleting an irrelevant row the model should retain its original label.)

Probing by altering or deleting relevant rows requires human annotation of relevant rows for each table-hypothesis pair. We used MTurk to annotate the relevance of rows in the development and the test sets, with turkers identifying the relevant rows for each table-hypothesis pair.

Inter-annotator Agreement. We employed majority voting to derive ground truth labels from multiple annotations for each row. The inter-annotator agreement macro F1 score for each of the four datasets is over 90% and the average Fleiss’ kappa is 78 (std: 0.22). This suggests good inter-annotator agreement. In 82.4% of cases, at least 3 out of 5 annotators marked the same relevant rows.

Results and Analysis We examined the response of the model when relevant rows are deleted. Figure 4 shows the label transitions. The fact that even after the deletion of relevant rows, **ENTAIL** and **CONTRADICT** predictions don’t change to **NEUTRAL** a large percentage of times (mostly the original label remains unchanged and at other

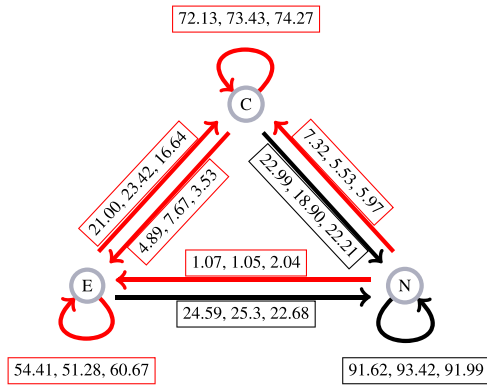


Figure 4: Changes in model predictions after deletion of relevant rows. Red lines represent invalid transitions while black lines represent valid transitions. The directed edges are labeled in the same manner as they are in Figure 3.

Dataset	α_1	α_2	α_3	Average
ENTAIL	75.41	74.70	77.31	75.80
NEUTRAL	8.39	6.58	8.01	7.66
CONTRADICT	77.02	81.10	77.80	78.64
Average	53.60	54.14	54.35	

Table 6: Percentage of invalid transitions following deletion of relevant rows. For an ideal model, all these numbers should be zero.

times, it changes incorrectly), indicates that the model is likely utilizing spurious statistical patterns in the data for making the prediction.

We summarize the combined invalid transitions for each label in Table 6. We see that the percentage of invalid transitions is considerably higher compared to random row deletion in Figure 3.⁸ The large percentage of invalid transitions in the ENTAIL and CONTRADICT cases indicates a rather high utilization of spurious statistical patterns by the model to arrive at its answers.

5.3 Human vs Model Evidence Selection

We further analyze the model’s capability for selecting relevant evidence by comparing it with human annotators. All rows that alter the model predictions during automatic row deletion are considered as *model relevant rows* and are compared to the human-annotated relevant rows. We only consider the subset of 4,600 (from 7,200 annotated dev/test sets pairs) hypothesis-table pairs

⁸Note that the dashed black lines from Figure 3 are now red in Figure 4, indicating invalid transitions.

with ENTAIL and CONTRADICT labels, where deleting a relevant row should change the prediction to NEUTRAL.⁹

Results and Analysis On the human-annotated relevant rows, the model has an average precision of 41.0% and a recall of 40.9%. Further analysis reveals that the model (a) uses all relevant rows in 27% cases, (b) uses incorrect or no rows as evidence in 52% of occurrences, and (c) is only partially accurate in identifying relevant rows in the remaining 21% of examples. Upon further analyzing the cases in (b), we observed that the model actually ignores premises completely in 88% (of 52%) of cases. This accounts for 46% (absolute) of all occurrences. In comparison, in the human-annotated data, such cases only amount to < 2%.

Although, the model’s predictions are 70% correct in the 4,600 examples, only 21% can be attributed to using all relevant evidence. The correct label in 37% of the 4,600 examples is from irrelevant rows, with the remaining 12% of correct predictions use some, but not all, relevant rows.

We can conclude from the findings in this section that the model does not seem to need all the relevant evidence to arrive at its predictions, raising questions about trust in its predictions.

6 Probing with Counterfactual Examples

Since INFO TABS is a dataset of facts based on Wikipedia, pre-trained language models such as RoBERTa, trained on Wikipedia and other publicly available text, may have already encountered information in INFO TABS during pre-training. As a result, NLI models built on top of RoBERTa_L can learn to infer a hypothesis using the knowledge of the pre-trained language model. More specifically, the model may be relying on “*confirmation bias*”, in which it selects evidence/patterns from both premise and hypothesis that matches its prior knowledge. While world knowledge is necessary for table NLI (Neeraja et al., 2021), models should still treat the premise as the primary evidence.

Counterfactual examples can help test whether the model is grounding its inference on the evidence provided in the tabular premise. In such examples, the tabular premise is modified such that the content does not reflect the real world. In

⁹We did not include the 2,400 NEUTRAL examples pairs and the ambiguous 200 ENTAIL or CONTRADICT examples that had no relevant rows as per the consensus annotation.

Breakfast in America	
Released	29 March 1979
Recorded	May–December 1978
Studio	The Village Recorder (Studio B) in Los Angeles
Genre	pop ; art rock ; soft rock
Length	125 minutes
Label	A&M
Producer	Peter Henderson, Supertramp

$\widehat{H}5$: **Breakfast in America** has a running time of over 3 hrs.

Figure 5: Counterfactual table-hypothesis pair created from Figure 1 and Figure 2. Only the values of ‘Length’ rows are swapped, rest of the rows from Figure 1 are copied over.

this study, we limit ourselves to modifying only the **ENTAIL** and **CONTRADICT** examples. We omit the **NEUTRAL** cases because the majority of them in INFO TABS involve out-of-table information; producing counterfactuals for them is much harder and involves the laborious creation of new rows with the right information.

The task of creating counterfactual tables presents two challenges. First, the modified tables should not be self-contradictory. Second, we need to determine the labels of the associated hypotheses after the table is modified. We employ a simple approach to generate counterfactuals that addresses both challenges. We use the evidence selection data (§5.2) to gather all premise-hypothesis pairs that share relevant keys such as ‘Born’, ‘Occupation’, and so forth. Counterfactual tables are generated by swapping the values of relevant keys from one table to another.¹⁰

Figure 5 shows an example. We create counterfactuals from the premises in Figure 1 and Figure 2 by swapping their **Length** rows. We also swap the hypotheses ($\widehat{H}1$ and $\widehat{H}5$) aligned to the **Length** rows in both premises by replacing the title expression **Bridesmaids** in $\widehat{H}5$ with **Breakfast in America** and *vice versa*. The simple procedure en-

¹⁰There may still be a few cases of self-contradiction, but we expect that such invalid cases would not exist in the rows that are relevant to the hypothesis.

Bridesmaids	
Released	29 March 1979
Recorded	May–December 1978
Studio	The Village Recorder (Studio B) in Los Angeles
Genre	pop ; art rock ; soft rock
Length	46:06
Label	A&M
Producer	Peter Henderson, Supertramp

$\widehat{H}1$: **Bridesmaids** is a pop album with a length of 46 minutes.

$\widehat{H}2$: **Bridesmaids** was released at the end of 1979.

$\widehat{H}3$: Most of **Bridesmaids** was recorded in the last month of 1978.

$\widehat{H}4$: **Bridesmaids** has 6 tracks.

Figure 6: A counterfactual tabular premise and the associated hypotheses created from Figures 1 and 2. The hypotheses $\widehat{H}1$ is entailed by the premise, $\widehat{H}2$ contradicts it, and $\widehat{H}3$ and $\widehat{H}4$ are neutral.

ures that the hypotheses labels are left unchanged in the process, resulting in high-quality data.

In addition, we also generated counterfactuals by swapping the table title and associated expressions in the hypotheses with the title of another table, resulting in a counterfactual table-hypothesis pair, as in the row swapping strategy. Figure 6 shows an example created from the premises in Figure 1 and Figure 2 by swapping the title rows **Breakfast in America** and **Bridesmaids**. The title expression in all hypotheses in Figure 1 are also replaced by **Bridesmaids**. This strategy also preserves the hypothesis label similar to row swapping.

The above approaches are *Label Preserving* as they do not alter the entailment labels. Counterfactual pairs with flipped labels are important for filtering out the contribution of artifacts or other spurious correlations that originate from a hypothesis (see §4). So, in addition, we also created counterfactual table-hypothesis pairs where the original labels are flipped. These counterfactual cases are, however, non-trivial to generate automatically, and are therefore created manually. To create the *Label-Flipped* counterfactual data, three annotators manually modified tables from

Model	Original	Label	Label
	$mean_{(stdev)}$	Preserved	Flipped
<i>Train Set (without NEUTRAL)</i>			
Prem+Hypo	94.38 _(0.39)	78.53 _(0.65)	48.70 _(0.72)
Hypo-Only	99.94 _(0.06)	82.23 _(0.65)	00.06 _(0.01)
α_1 Set (without NEUTRAL)			
Prem+Hypo	71.99 _(0.69)	69.65 _(0.78)	44.01 _(0.72)
Hypo-Only	60.89 _(0.76)	58.19 _(0.91)	27.68 _(0.65)

Table 7: Results of the Hypothesis-only and Prem+Hypo models on the gold and counterfactual examples.

the *Train* and α_1 datasets corresponding to **ENTAIL** and **CONTRADICT** labels, producing 885 counterfactual examples from the *Train* set and 942 from the α_1 set. The annotators cross-checked the labels to determine annotation accuracy, which was 88.45% for the *Train* set and 86.57% for the α_1 set.

Results and Analysis We tested both hypothesis-only and full (Prem+Hypo) models on the counterfactual examples created above, without fine-tuning on a subset of these examples. The results are presented in Table 7 where each cell represents average accuracy and standard deviation (subscript) over 100 sets of 80% randomly sampled counterfactual examples. We see that the (Prem+Hypo) model is not robust to counterfactual perturbations. On the label-flipped counterfactuals, the performance drops down to close to a random prediction (48.70% for *Train* and 44.01% for α_1). The performance on the label-preserved counterfactuals is relatively better which leads us to conjecture that the model largely exploits artifacts in hypotheses.

Due to over-fitting, the *Train* set has a larger drop of 15.85%, compared to only 2.70% on the α_1 set on the label-preserved examples. Moreover, the drop in performance for both *Prem+Hypo* and *Hypo-Only* models is comparable to their performance drop on the original table-hypothesis pairs. This shows that, regardless of whether the relevant information in the premise is accurate, both models rely substantially on hypothesis artifacts. On the *Label-Flipped* counterfactuals, the large drop in accuracy could be due to both ambiguous hypothesis artifacts or counterfactual information.

To disentangle these two factors, we can take advantage of the fact that the counterfactual examples are constructed from, and hence paired

Prem+Hypo		Hypo-Only	Dataset	
C-THP	O-THP	O-Hypo	<i>Train</i>	α_1
✓	✗	✗	0.00	11.43
✗	✓	✗	0.00	11.79
✓	✗	✓	3.57	6.48
✗	✓	✓	49.36	33.12

Table 8: Performance of the full and hypothesis-only models on the original and counterfactual examples. O-THP and C-THP represent original and counterfactual table-hypothesis pairs; O-Hypo represents hypotheses from the original data; ✓ represents correct predictions and ✗ represents incorrect predictions.

with, the original examples. This allows us to examine pairs of examples where the full model makes an incorrect prediction on one, but not the other. Especially of interest are the cases where the full model makes a correct prediction on the original example, but not on the corresponding counterfactual example.

Table 8 shows the results of this analysis. Each row represents a condition corresponding to whether the full and the hypothesis-only models are correct on the original example. The two cases of interest, described above, correspond to the second and fourth rows of the table. The second row shows the case where the full model is correct on the original example (and not on the counter-factual example), but the hypothesis-only model is not. Since we can discount the impact of hypothesis bias in these examples, the error in the counter-factual version could be attributed to reliance on pre-trained knowledge. Unsurprisingly, there are no such examples in the training set. In the α_1 set, we see a substantial fraction of counterfactual examples (11.79%) belong to this category. The last row considers the case where the hypothesis-only model is correct. We see that this accounts for a larger fraction of the counterfactual errors, both in the training and the α_1 sets. Among these examples, *despite* the (albeit unfortunate) fact that the hypothesis alone can support a correct prediction, the model’s reliance on its pre-trained knowledge leads to errors in the counterfactual cases.

The results, taken in aggregate, suggest that the model produces predictions based on hypothesis artifacts and pre-trained knowledge rather than

the evidence presented to it, thus impacting its robustness and generalization.

7 Inoculation by Fine-Tuning

Our probing experiments demonstrate that the models, trained on the INFO TABS training set, failed along all three dimensions that we investigated. This leads us to the following question: *Can additional fine-tuning with perturbed examples help?*

Liu et al. (2019a) point out that poor performance on challenging datasets can be ascribed to either a weakness in the model, a lack of diversity in the dataset used for training, or information leakage in the form of artifacts.¹¹ They suggest that models can be further fine-tuned on a few challenging examples to determine the possible source of degradation. Inoculation can lead to one of three outcomes: (a) **Outcome 1:** The performance gap between the challenge and the original test sets reduces, possibly due to addition of diverse examples, (b) **Outcome 2:** Performance on both the test sets remains unchanged, possibly because of the model’s inability to adapt to the new phenomena or the changed data distribution, or, (c) **Outcome 3:** Performance degrades on the test set, but improves on the challenge set, suggesting that adding new examples introduces ambiguity or contradictions.

We conducted two sets of inoculation experiments to help categorize performance degradation of our models into one of these three categories. For each experiment described below, we generated additional inoculation datasets with 100, 200, and 300 examples to inoculate the original task-specific RoBERTa_L models trained on both premises and hypotheses. As in the original inoculation work, we created these adversarial datasets by sub-sampling inclusively, i.e., the smaller datasets are subsets of the larger ones. Following the training protocol in Liu et al. (2019a), we tried learning rates of 10^{-4} , 5×10^{-5} , and 10^{-5} . We performed inoculation for a maximum of 30 epochs with early stopping based on the development set accuracy. We found that with the first two learning rates, the model does not converge, and underperforms on the development set. The model performance was best with the learning rate of 10^{-5} , which we used throughout the inoculation experiments. The standard deviation

¹¹Model weakness is the inherent inability of a model (or a model family) to handle certain linguistic phenomena.

#Samples	α_1	α_2	α_3
0 (w/o Ino)	74.88	65.55	64.94
100	67.44	62.17	58.51
200	67.34	61.88	58.61
300	67.24	61.84	58.62

Table 9: Performance of the inoculated models on the original INFO TABS test sets.

#Samples	Original	Label Preserved	Label Flipped
<i>Train Set (w/o NEUTRAL)</i>			
0 (w/o Ino)	99.44	92.98	53.92
100	97.24	95.58	79.25
200	97.24	95.65	78.75
300	97.24	95.64	78.74
α_1 Set (w/o NEUTRAL)			
0 (w/o Ino)	68.94	69.56	51.48
100	68.05	65.67	57.91
200	68.37	66.29	57.49
300	68.36	66.29	57.49

Table 10: Performance of the inoculated models on the hypothesis perturbed INFO TABS sets.

over 100 sample splits for all experiments was ≤ 0.91 .

Annotation Artifacts Table 9 shows the performance of the inoculated models on the original INFO TABS test sets, and Table 10 shows the results on the hypothesis-perturbed examples (from §4). We see that fine-tuning on the hypothesis-perturbed examples decreases performance on the original α_1 , α_2 , and α_3 test sets, but performance improves on the more difficult label-flipped examples of the hypothesis-perturbed test set.

Counterfactual Examples Tables 11 and 12 show the performance of models inoculated on the original INFO TABS test sets and the counterfactual examples from §6, respectively. Once again, we see that fine-tuning on counterfactual examples improves performance on the adversarial counterfactual examples test set, at the cost of performance on the original test sets.

Analysis We see that both experiments above belong to Outcome 3, where the performance improves on the challenge set, but degrades on the test set(s). The change in the distribution of inputs hurts the model: we conjecture that this may be because the RoBERTa_L model exploits data artifacts in the original dataset but fails to do so for the challenge dataset and vice versa.

#Samples	α_1	α_2	α_3
0 (w/o Ino)	74.88	65.55	64.94
100	69.72	63.88	59.66
200	69.88	63.78	58.89
300	67.34	62.23	57.58

Table 11: Performance after inoculation by fine-tuning on the original INFO TABS test sets.

#Samples	Original	Label Preserved	Label Flipped
<i>Train Set (w/o NEUTRAL)</i>			
0 (w/o Ino)	94.38	78.53	48.70
100	91.82	84.61	57.62
200	92.46	84.92	59.43
300	91.08	83.54	63.58
α_1 Set (w/o NEUTRAL)			
0 (w/o Ino)	71.99	69.65	44.01
100	66.05	75.03	50.40
200	65.86	75.03	50.57
300	65.59	74.23	52.09

Table 12: Performance after inoculation fine-tuning on the INFO TABS counterfactual example sets.

We expect our model to handle both original and challenge datasets, at least after fine-tuning (i.e., it should belong to Outcome 1). Its failure points to the need for better models or training regimes.

8 Discussion and Related Work

What Did We Learn? Firstly, through systematic probing, we have shown that despite good performance on the evaluation sets, the model for tabular NLI fails at reasoning. From the analysis of hypothesis perturbations (§4), we show that the model heavily relies on correlations between a hypothesis’ sentence structure and its label. Models should be systematically evaluated on adversarial sets like α_2 for robustness and sensitivity. This observation is concordant with multiple studies that probe deep learning models on adversarial examples in a variety of tasks such as question answering, sentiment analysis, document classification, natural language inference, and so forth. (e.g., Ribeiro et al., 2020; Richardson et al., 2020; Goel et al., 2021; Lewis et al., 2021; Tarunesh et al., 2021).

Secondly, the model does not look at correct evidence required for reasoning, as is evident from the evidence-selection probing (§5). Rather, it leverages spurious patterns and statistical cor-

relations to make predictions. A recent study by Lewis et al. (2021) on question-answering shows that models indeed leverage spurious patterns to answer a large fraction (60–70%) of questions.

Thirdly, from counterfactual probes (§6), we found that the model relies on knowledge of pre-trained language models than on tabular evidence as the primary source of knowledge for making predictions. This is in addition to the spurious patterns or hypothesis artifacts leveraged by the model. Similar observations are made by Clark and Etzioni (2016), Jia and Liang (2017), Kaushik et al. (2020), Huang et al. (2020), Gardner et al. (2020), Tu et al. (2020), Liu et al. (2021), Zhang et al. (2021), and Wang et al. (2021b) for unstructured text.

Finally, from the inoculation study (§7), we found that fine-tuning on challenge sets improves model performance on challenge sets but degrades on the original α_1 , α_2 , and α_3 test sets. That is, changes in the data distribution during training have a negative impact on model performance. This adds weight to the argument that the model relies excessively on data artifacts.

Benefit of Tabular Data Unlike unstructured data, where creating challenge datasets may be more difficult (e.g., Ribeiro et al., 2020; Goel et al., 2021; Mishra et al., 2021), we can analyze semi-structured data more effectively. Although connected with the title, the rows in the table are still independent, linguistically and otherwise. Thus, controlled experiments are easier to design and study. For example, the analysis done for evidence selection via multiple table perturbation operations such as row deletion and insertion is possible mainly due to the tabular nature of the data. Such granularity and component-independence is generally absent for raw text at the token, sentence, and even paragraph level. As a result, designing suitable probes with sufficient coverage can be a challenging task, and can require more manual effort.

Additionally, probes defined on one tabular dataset (INFO TABS in our case) can be easily ported to other tabular datasets such as WikiTableQA (Pasupat and Liang, 2015), TabFact (Chen et al., 2020b), HybridQA (Chen et al., 2020c; Zayats et al., 2021; Oguz et al., 2020), OpenTableQA (Chen et al., 2021), ToTTo (Parikh et al., 2020), Turing Tables (Yoran et al., 2021), and LogicTable (Chen et al., 2020a). Moreover, such probes

can be used to better understand the behavior of various tabular reasoning models (e.g., Müller et al., 2021; Herzig et al., 2020; Yin et al., 2020; Iida et al., 2021; Pramanick and Bhattacharya, 2021; Glass et al., 2021; and others).

Interpretability for NLI Models For classification tasks such as NLI, correct predictions do not always mean that the underlying model is employing correct reasoning. More work is needed to make models interpretable, either through explanations or by pointing to the evidence that is used for predictions (e.g., Feng et al., 2018; Serrano and Smith, 2019; Jain and Wallace, 2019; Wiegrefe and Pinter, 2019; DeYoung et al., 2020; Paranjape et al., 2020; Hewitt and Liang, 2019; Richardson and Sabharwal, 2020; Niven and Kao, 2019; Ravichander et al., 2021). Many recent shared tasks on reasoning over semi-structured tabular data (such as SemEval 2021 Task 9 [Wang et al., 2021a] and FEVEROUS [Aly et al., 2021]) have highlighted the importance of, and the challenges associated with, evidence extraction for claim verification.

Finally, NLI models should be tested on multiple test sets in adversarial settings (e.g., Ribeiro et al., 2016, 2018a,b; Alzantot et al., 2018; Iyyer et al., 2018; Glockner et al., 2018; Naik et al., 2018; McCoy et al., 2019; Nie et al., 2019; Liu et al., 2019a) focusing on particular properties or aspects of reasoning, such as perturbed premises for evidence selection, zero-shot transfer (α_3), counterfactual premises or alternate facts, and contrasting hypotheses via perturbation (α_2). Such behavioral probing by evaluating on multiple test-only benchmarks and controlled probes is essential to better understand both the abilities and the weaknesses of pre-trained language models.

9 Conclusion

This paper presented a targeted probing study to highlight the limitations of tabular inference models using a case study on a tabular NLI task on INFO TABS. Our findings show that despite good performance on standard splits, a RoBERTa-based tabular NLI model, fine-tuned on the existing pre-trained language model, fails to select the correct evidence, makes incorrect predictions on adversarial hypotheses, and is not grounded in provided evidence—counterfactual or otherwise. We expect that insights from the study can help in designing rationale selection techniques based on

structural constraints for tabular inference and other tasks. While inoculation experiments showed partial success, diverse data augmentation may help mitigate challenges. However, annotation of such data can be expensive. It may also be possible to train models to satisfy domain-based constraints (e.g., Li et al., 2020) to improve model robustness. Finally, probing techniques described here may be adapted to other NLP tasks involving tables such as tabular question answering and tabular text generation.

Acknowledgments

We thank the reviewing team for their valuable feedback. This work is partially supported by NSF grants #1801446 and #1822877.

References

- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. 2021. The fact extraction and VERification over unstructured and structured information (FEVEROUS) shared task. In *Proceedings of the Fourth Workshop on Fact Extraction and VERification (FEVER)*, pages 1–13, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.fever-1.1>
- Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. Generating natural language adversarial examples. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2890–2896, Brussels, Belgium. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D18-1316>
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D15-1075>
- Wenhu Chen, Ming-Wei Chang, Eva Schlinger, William Yang Wang, and William W. Cohen.

2021. Open question answering over tables and text. In *International Conference on Learning Representations*.
- Wenhu Chen, Jianshu Chen, Yu Su, Zhiyu Chen, and William Yang Wang. 2020a. Logical natural language generation from open-domain tables. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7929–7942, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.708>
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2020b. Tabfact: A large-scale dataset for table-based fact verification. In *International Conference on Learning Representations*.
- Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Yang Wang. 2020c. HybridQA: A dataset of multi-hop question answering over tabular and textual data. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1026–1036, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.91>
- Peter Clark and Oren Etzioni. 2016. My computer is an honor student—but how intelligent is it? Standardized tests as a measure of AI. *AI Magazine*, 37(1):5–12. <https://doi.org/10.1609/aimag.v37i1.2636>
- Ido Dagan, Dan Roth, Mark Sammons, and Fabio Massimo Zanzotto. 2013. Recognizing textual entailment: Models and applications. *Synthesis Lectures on Human Language Technologies*, 6(4):1–220. <https://doi.org/10.2200/S00509ED1V01Y201305HLT023>
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jay DeYoung, Sarthak Jain, Nazneen Fatema Rajani, Eric Lehman, Caiming Xiong, Richard Socher, and Byron C. Wallace. 2020. ERASER: A benchmark to evaluate rationalized NLP models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4443–4458, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.408>
- Julian Eisenschlos, Syrine Krichene, and Thomas Müller. 2020. Understanding tables with intermediate pre-training. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 281–296, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.27>
- 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, Brussels, Belgium. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D18-1407>
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khashabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. Evaluating models’ local decision boundaries via contrast sets. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.findings-emnlp.117>
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets. 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 1161–1166, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1107>

- Michael Glass, Mustafa Caim, Alfio Gliozzo, Saneem Chemmengath, Vishwajeet Kumar, Rishav Chakravarti, Avi Sil, Feifei Pan, Samarth Bharadwaj, and Nicolas Rodolfo Fauceglia. 2021. Capturing row and column semantics in transformer based question answering over tables. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1212–1224, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.96>
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. Breaking NLI systems with sentences that require simple lexical inferences. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-2103>
- Karan Goel, Nazneen Fatema Rajani, Jesse Vig, Zachary Taschdjian, Mohit Bansal, and Christopher Ré. 2021. Robustness gym: Unifying the NLP evaluation landscape. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Demonstrations*, pages 42–55, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-demos.6>
- Vivek Gupta, Maitrey Mehta, Pegah Nokhiz, and Vivek Srikumar. 2020. INFOTABS: Inference on tables as semi-structured data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2309–2324, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.210>
- 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. <https://doi.org/10.18653/v1/N18-1170>
- Jonathan Herzig, Pawel Krzysztof Nowak, Thomas Müller, Francesco Piccinno, and Julian Eisenschlos. 2020. TaPas: Weakly supervised table parsing via pre-training. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4320–4333, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.398>
- John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. 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 2733–2743, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1275>
- William Huang, Haokun Liu, and Samuel R. Bowman. 2020. Counterfactually-augmented SNLI training data does not yield better generalization than unaugmented data. In *Proceedings of the First Workshop on Insights from Negative Results in NLP*, pages 82–87, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.insights-1.13>
- Hiroshi Iida, Dung Thai, Varun Manjunatha, and Mohit Iyyer. 2021. TABBIE: Pretrained representations of tabular data. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3446–3456, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.270>
- Mohit Iyyer, John Wieting, Kevin Gimpel, and Luke Zettlemoyer. 2018. Adversarial example generation with syntactically controlled paraphrase networks. 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 1875–1885, New Orleans, Louisiana. Association for Computational Linguistics. <https://doi.org/10.18653/v1/N18-1170>

- 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, Minneapolis, Minnesota. 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. <https://doi.org/10.18653/v1/D17-1215>
- Divyansh Kaushik, Eduard Hovy, and Zachary Lipton. 2020. Learning the difference that makes a difference with counterfactually-augmented data. In *International Conference on Learning Representations*.
- Patrick Lewis, Pontus Stenetorp, and Sebastian Riedel. 2021. Question and answer test-train overlap in open-domain question answering datasets. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1000–1008, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.86>
- Tao Li, Parth Anand Jawale, Martha Palmer, and Vivek Srikumar. 2020. Structured tuning for semantic role labeling. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8402–8412, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.744>
- 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. *arXiv preprint arXiv:1907.11692. Version 1*.
- Zeyu Liu, Yizhong Wang, Jungo Kasai, Hannaneh Hajishirzi, and Noah A. Smith. 2021. Probing across time: What does RoBERTa know and when? In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 820–842, Punta Cana, Dominican Republic. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.findings-emnlp.71>
- 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. <https://doi.org/10.18653/v1/P19-1334>
- Anshuman Mishra, Dhruvesh Patel, Aparna Vijayakumar, Xiang Lorraine Li, Pavan Kapanipathi, and Kartik Talamadupula. 2021. Looking beyond sentence-level natural language inference for question answering and text summarization. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1322–1336, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.104>
- Thomas Müller, Julian Eisenschlos, and Syrine Krichene. 2021. TAPAS at SemEval-2021 task 9: Reasoning over tables with intermediate pre-training. In *Proceedings of the 15th International Workshop on Semantic Evaluation (SemEval-2021)*, pages 423–430, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.semeval-1.51>
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

- J. Neeraja, Vivek Gupta, and Vivek Srikumar. 2021. Incorporating external knowledge to enhance tabular reasoning. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2799–2809, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.224>
- Yixin Nie, Yicheng Wang, and Mohit Bansal. 2019. Analyzing compositionality-sensitivity of NLI models. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):6867–6874. <https://doi.org/10.1609/aaai.v33i01.33016867>
- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1459>
- Barlas Oguz, Xilun Chen, Vladimir Karpukhin, Stan Peshterliev, Dmytro Okhonko, Michael Schlichtkrull, Sonal Gupta, Yashar Mehdad, and Scott Yih. 2020. Unified open-domain question answering with structured and unstructured knowledge. *arXiv preprint arXiv:2012.14610, Version 2*.
- Bhargavi Paranjape, Mandar Joshi, John Thickstun, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2020. An information bottleneck approach for controlling conciseness in rationale extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1938–1952, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.153>
- Ankur Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. 2020. ToTTo: A controlled table-to-text generation dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1173–1186, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.emnlp-main.89>
- Panupong Pasupat and Percy Liang. 2015. Compositional semantic parsing on semi-structured tables. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1470–1480, Beijing, China. Association for Computational Linguistics. <https://doi.org/10.3115/v1/P15-1142>
- 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. <https://doi.org/10.18653/v1/S18-2023>
- Aniket Pramanick and Indrajit Bhattacharya. 2021. Joint learning of representations for webtables, entities and types using graph convolutional network. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1197–1206, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.102>
- Abhilasha Ravichander, Yonatan Belinkov, and Eduard Hovy. 2021. Probing the probing paradigm: Does probing accuracy entail task relevance? In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3363–3377, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.295>
- 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, KDD ’16*, pages 1135–1144, New York, NY, USA. Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939778>
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018a. Anchors: High-precision model-

- agnostic explanations. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018b. Semantically equivalent adversarial rules for debugging NLP models. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865, Melbourne, Australia. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P18-1079>
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. Beyond accuracy: Behavioral testing of NLP models with CheckList. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2020.acl-main.442>
- Kyle Richardson, Hai Hu, Lawrence Moss, and Ashish Sabharwal. 2020. Probing natural language inference models through semantic fragments. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8713–8721. <https://doi.org/10.1609/aaai.v34i05.6397>
- Kyle Richardson and Ashish Sabharwal. 2020. What does my QA model know? Devising controlled probes using expert knowledge. *Transactions of the Association for Computational Linguistics*, 8:572–588. https://doi.org/10.1162/tacl_a_00331
- 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, Florence, Italy. Association for Computational Linguistics. <https://doi.org/10.18653/v1/P19-1282>
- Ishan Tarunesh, Somak Aditya, and Monojit Choudhury. 2021. Trusting RoBERTa over BERT: Insights from checklisting the natural language inference task. *arXiv preprint arXiv:2107.07229. Version 1*.
- Lifu Tu, Garima Lalwani, Spandana Gella, and He He. 2020. An empirical study on robustness to spurious correlations using pre-trained language models. *Transactions of the Association for Computational Linguistics*, 8:621–633. https://doi.org/10.1162/tacl_a_00335
- Nancy Xin Ru Wang, Diwakar Mahajan, Marina Danilevsky, and Sara Rosenthal. 2021a. Semeval-2021 task 9: Fact verification and evidence finding for tabular data in scientific documents (sem-tab-facts). In *SemEval@ACL/IJCNLP*, pages 317–326.
- Siyuan Wang, Wanjun Zhong, Duyu Tang, Zhongyu Wei, Zhihao Fan, Daxin Jiang, Ming Zhou, and Nan Duan. 2021b. Logic-driven context extension and data augmentation for logical reasoning of text. *arXiv preprint arXiv:2105.03659. Version 1*.
- Sarah Wiegrefe and Yuval Pinter. 2019. Attention is not explanation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 11–20, Hong Kong, China. Association for Computational Linguistics. <https://doi.org/10.18653/v1/D19-1002>
- Pengcheng Yin, Graham Neubig, Wen-tau Yih, and Sebastian Riedel. 2020. TaBERT: Pretraining for joint understanding of textual and tabular data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8413–8426, Online. Association for Computational Linguistics.
- Ori Yoran, Alon Talmor, and Jonathan Berant. 2021. Turning tables: Generating examples from semi-structured tables for endowing language models with reasoning skills. *arXiv preprint arXiv:2107.07261. Version 1*.
- Vicky Zayats, Kristina Toutanova, and Mari Ostendorf. 2021. Representations for question answering from documents with tables and text. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2895–2906, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.eacl-main.253>
- Chong Zhang, Jieyu Zhao, Huan Zhang, Kai-Wei Chang, and Cho-Jui Hsieh. 2021. Double perturbation: On the robustness of robustness and

counterfactual bias evaluation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3899–3916, Online. Association for Computational Linguistics. <https://doi.org/10.18653/v1/2021.naacl-main.305>

Appendix

In Section 5.1, we defined four types of row-agnostic table modifications: (a) row deletion, (b) row insertion, (c) row-value update, and (d) row permutation and presented the first one there. We present details of the rest here along with the respective impact on the α_1 , α_2 , and α_3 test sets.

Row Insertion. When we insert new information that does not contradict an existing table,¹² original predictions should be retained in almost all cases. Very rarely, **NEUTRAL** labels may change to **ENTAIL** or **CONTRADICT**. For example, adding the *Singles* row below to our running example table doesn't change labels for any hypothesis except the H4 label (see Figure 1) changing to **CONTRADICT** with the additional information.

Singles The Logical Song; Breakfast in America; Goodbye Stranger; Take the Long Way Home

Figure 7 shows the possible label changes after new row insertion as a directed labeled graph, and the results are summarized in Table 13. Note that all transitions from **NEUTRAL** are valid upon row insertion, although not all may be accurate.

Row Update. In case of row update, we only change a portion of a row value. Whole row value substitutions are examined separately as composite operations of deletion followed by insertion. Unlike a whole row update, changing only a portion of a row is non-trivial. We must ensure that the updated value is appropriate for the key in question and also avoid self-contradictions. To satisfy these constraints, we update a row with a value from a random table with the same key and only update values in multi-valued rows. A row update operation may have an effect on all labels.

¹²To ensure that the information added is not contradictory to existing rows, we only add rows with new keys instead of changing values for the existing keys.

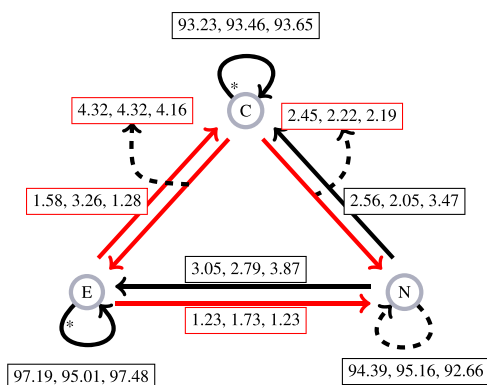


Figure 7: Changes in model predictions after new row insertion. (Notation similar to Figure 3).

Dataset	α_1	α_2	α_3	Average
ENTAIL	2.81	4.99	2.51	3.44
NEUTRAL	0	0	0	0
CONTRADICT	6.77	6.54	6.35	6.55
Average	3.19	3.84	2.95	—

Table 13: Percentage of invalid transitions after new row insertion. For an ideal model, all these numbers should be zero.

Though feasible, we consider the transitions from **CONTRADICT** to **ENTAIL** to be prohibited. Unlike **ENTAIL** to **CONTRADICT** transitions, these transitions would be extremely rare as values are updated randomly, regardless of their semantics. For example, if we substitute *pop* in the multi-valued key *Genre* in our running example with another genre, the hypothesis H1 is likely to change to **CONTRADICT**.

Since we are updating a single value from a multi-valued key, the changes to the table are minimal and may not be perceived by the model. As a result, we should expect row updates to have lower impact on model predictions. This appears to be the case, as evidenced by the results in Figure 8, which show that the labels do not change drastically after update. The results in Figure 8 are summarized in Table 14.

Row Permutation. By design of the premises, the order of their rows should have no effect on hypotheses labels. In other words, the labels should be invariant to row permutation. However, from Figure 9, it is evident that even a simple shuffling of rows, where no information has been tampered with, can have a notable effect on performance.

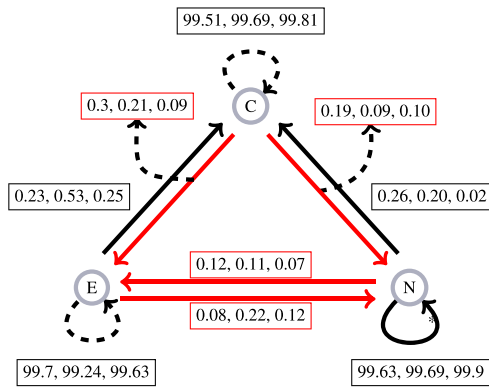


Figure 8: Changes in model predictions after row value update. (Notation similar to Figure 3).

Dataset	α_1	α_2	α_3	Average
ENTAIL	0.08	0.22	0.12	0.14
NEUTRAL	0.12	0.11	0.09	0.11
CONTRADICT	0.49	0.30	0.19	0.33
Average	0.23	0.21	0.13	–

Table 14: Percentage of invalid transitions after row value update. For an ideal model, all these numbers should be zero.

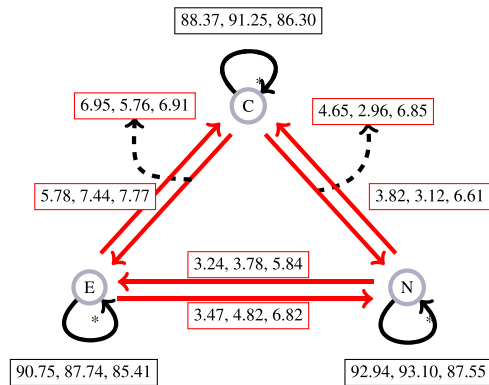


Figure 9: Changes in model predictions after shuffling of table rows. (Notation similar to Figure 3.)

This shows that the model is relying on row positions incorrectly, while the semantics of a table is order invariant. We summarize the combined invalid transitions from Figure 9 in Table 15.

Irrelevant Row Deletion. Ideally, deletion of an irrelevant row should have no effect on a hypothesis label. The results in Figure 10 and in Table 16 show that even irrelevant rows have an effect on model predictions. This further illustrates that the seemingly accurate model predictions are not appropriately grounded on evidence.

Dataset	α_1	α_2	α_3	Average
ENTAIL	9.25	12.2	14.6	12.02
NEUTRAL	7.1	6.8	12.5	8.79
CONTRADICT	11.6	8.76	13.7	11.36
Average	9.34	9.26	13.6	–

Table 15: Percentage of invalid transitions after row permutations. For an ideal model, all these numbers should be zero.

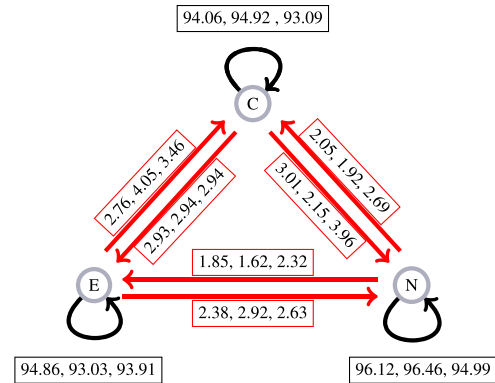


Figure 10: Change in model predictions after deletion of an irrelevant row. (Notation similar to Figure 3.)

Dataset	α_1	α_2	α_3	Average
ENTAIL	5.14	6.97	6.09	6.07
NEUTRAL	3.9	3.54	5.01	4.15
CONTRADICT	5.94	5.09	6.91	5.98
Average	4.99	5.2	6.01	–

Table 16: Percentage of invalid transitions after deletion of irrelevant rows. For an ideal model, all these numbers should be zero.

Composition of Perturbation Operations In addition to probing individual operations, we can also study their compositions. For example, we could delete a row, and insert a different row, and so on. The composition of these operations have interesting properties with respect to the allowed transitions. For example, when an operation is composed with itself (e.g., two deletions), the set of valid label changes is the same as for the operation. A particularly interesting composition is deletion followed by an insertion, since this can be viewed as a row update. In Figure 11, we show the transition graph for the composition operation of row deletion followed by insertion and the summary of the possible transitions is presented in Table 17.

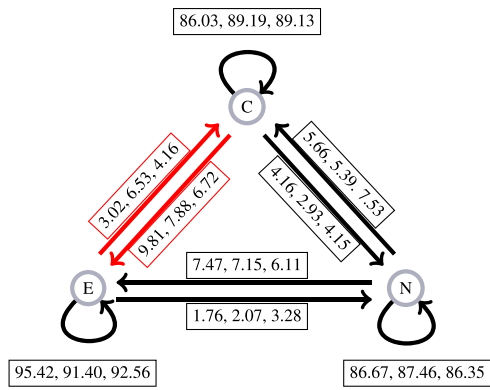


Figure 11: Changes in model predictions after deletion followed by an insert operation. (Notation similar to Figure 3.)

Dataset	α_1	α_2	α_3	Average
ENTAIL	3.02	6.53	4.16	4.57
NEUTRAL	0.00	0.00	0.00	0.00
CONTRADICT	9.81	7.88	6.71	8.13
Average	4.28	4.80	3.63	—

Table 17: Percentage of invalid transitions after deletion followed by an insertion operation. For an ideal model, all these numbers should be zero.