

LIREx: Augmenting Language Inference with Relevant Explanation

Xinyan Zhao,¹ V.G.Vinod Vydiswaran^{2,1}

¹School of Information; ²Department of Learning Health Sciences
University of Michigan, Ann Arbor, Michigan 48109 USA
{zhaoxy, vgvinodv}@umich.edu

Abstract

Natural language explanations (NLEs) are a special form of data annotation in which annotators identify rationales (most significant text tokens) when assigning labels to data instances, and write out explanations for the labels in natural language based on the rationales. NLEs have been shown to capture human reasoning better, but not as beneficial for natural language inference (NLI). In this paper, we analyze two primary flaws in the way NLEs are currently used to train explanation generators for language inference tasks. We find that the explanation generators do not take into account the variability inherent in human explanation of labels, and that the current explanation generation models generate spurious explanations. To overcome these limitations, we propose a novel framework, LIREx, that incorporates both a rationale-enabled explanation generator and an instance selector to select only relevant, plausible NLEs to augment NLI models. When evaluated on the standardized SNLI data set, LIREx achieved an accuracy of 91.87%, an improvement of 0.32 over the baseline and matching the best-reported performance on the data set. It also achieves significantly better performance than previous studies when transferred to the out-of-domain MultiNLI data set. Qualitative analysis shows that LIREx generates flexible, faithful, and relevant NLEs that allow the model to be more robust to spurious explanations. The code is available at <https://github.com/zhaoxy92/LIREx>.

Introduction

Natural language explanations (NLEs) provided at the time of assigning labels to data instances are likely to better capture human reasoning than the labels alone. NLEs are special forms of data annotation in which annotators identify both the class label and the rationales (most significant text tokens), and write out explanations in natural language. They have been suggested to potentially improve the performance and interpretability of deep learning-based models – i.e. either augmenting model performance by incorporating NLEs as additional contextual features, or explaining model decisions by training an explanation generator. Researchers have parsed NLEs into structured logical forms (Srivastava, Labutov, and Mitchell 2017; Hancock et al. 2018; Lee et al. 2020; Qin et al. 2020) or directly encoded them into a vector-based semantic representation (Fidler et al. 2017). Recent

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Premise: A man wearing a red uniform and helmet stands on his motorbike.
Hypothesis: A man sitting in a car.
Gold label: contradiction

Hypothesis: A man **sitting** in a car. **Explain:** A man cannot be both standing and sitting.
Hypothesis: A man sitting in a **car**. **Explain:** A motorbike is not a car.

(a) Example of different NL explanations using different rationales.

Premise: A man wearing a white shirt is playing the drums.
Hypothesis: A man is playing a musical instrument.
Gold label: entailment

Explain (entailment): Drums is a musical instrument.
Explain (neutral): Drumming does not always involve playing a musical instrument.
Explain (contradiction): A man playing the drums isn't playing a musical instrument.

(b) Example of correct and incorrect NL explanations generated for a same premise-hypothesis, but favoring different labels.

Figure 1: Examples of NLEs in NLI task.

success in language modeling and generation have enabled trained models to explicitly provide human-readable explanations for classification tasks (Kim et al. 2018; Huk Park et al. 2018; Camburu et al. 2018; Kumar and Talukdar 2020; Rajani et al. 2019). Similarly, studies such as (Rajani et al. 2019) have reported significant performance improvements on commonsense reasoning tasks by including NLEs in training a language generation model. However, these trends do not carry over to natural language inference (NLI) task in which a premise-hypothesis pair is expected to be classified into *entailment*, *neutral*, or *contradiction*. Previous studies on utilizing NLEs for NLI tasks have reported a drop in overall performance, even with powerful deep learning-based models such as LSTM (Camburu et al. 2018), RoBERTa, and GPT2 (Kumar and Talukdar 2020). We study this discrepancy in more detail and identify two primary issues, described below, with how NLEs have been incorporated for the NLI task so far.

Issue 1: Lack of rationale in NLE Generation Current approaches for explanation generation produce only one specific explanation for each data instance. However, these approaches ignore the variability in human reasoning and alternative explanations. Annotators could assign the same label to a data instance by considering different rationales. For example, given the premise and the hypothesis shown

in Figure 1a, it is easy to infer that the label should be *contradiction*. However, this label could be explained using two different rationales – indicated by “*sitting*” or by “*car*”. Ignoring this aspect would limit the application of NLE in NLI tasks because trustworthy explanations should be consistent with the appropriate rationale used by humans to interpret the label.

Issue 2: Inclusion of Spurious Explanations NLEs that are inconsistent with commonsense logic provide little help for model prediction (Camburu et al. 2020). For example, given the premise-hypothesis pair in Figure 1b, the explanation regarding the correct label (*entailment*) aligns with the fact that drums are indeed musical instruments. However, while tailored explanations could be generated to favor other labels – *neutral* and *contradiction*, such explanations are themselves factually incorrect. This happens because while deep learning-based text generators are powerful enough to generate readable sentences, they often lack commonsense reasoning ability (Zhou et al. 2020). When generating an explanation, such models are prone to output negating text if conditioned to the *contradiction* label, or output text with uncertainty if conditioned to the *neutral* label, without reasoning about the plausibility of the generated text.

Proposed Solution: Language Inference with Relevant Explanations To address the aforementioned issues observed in how NLEs have been incorporated in NLI models, we propose a novel framework for Language Inference with Relevant Explanations (LIREx). LIREx augments the NLI model with relevant, plausible NLEs produced and selected by a rationale-enabled explanation generator and an instance selector. We conduct detailed analysis to show not only that our model is able to bring significant performance improvement to the NLI task, but also that the generated explanations are highly aligned with human interpretation when evaluated on a relevance-based evaluation metric.

Related Work

NLEs have been studied along two main directions. The first direction focused on how explanations could be treated as contextual features to improve the model performance. Studies in this direction include (Li et al. 2016; Srivastava, Labutov, and Mitchell 2017; Wang et al. 2017; Hancock et al. 2018; Qin et al. 2020; Lee et al. 2020), in which the authors used semantic parsers to convert unstructured NLEs into structured feature-like logical forms. These logical forms could further benefit low-resource setting by weakly labeling more unlabeled data. One limitation of such approaches is that the localized contexts usually have limited ability to represent the semantic meaning of text and are often difficult to convert to logical forms when they get too complicated.

The second direction focused on training NLE generators to justify model predictions, usually as a post-hoc exercise. (Kim et al. 2018) trained textual explanation generators conditioned on the video frames and commands in self-driving cars to describe and justify the operated actions. (Huk Park et al. 2018) proposed an explanation module for

visual question answering and an activity recognition task, in which they first use encoders to jointly predict labels and infer the rationale regions of an image, and then generate text explanations by conditioning on the predicted labels and inferred rationales. (Camburu et al. 2018) suggested the inclusion of NLEs for NLI task by proposing e-SNLI, an expanded dataset that contains NLE annotations. They also jointly trained the prediction model and the explanation generation model conditioned on the predicted label. However, jointly training the two models led to a non-negligible loss in performance (by about 2 points in F1).

In recent studies, generated NLEs have been combined with original data for label prediction tasks. (Rajani et al. 2019) proposed CAGE for commonsense reasoning task, where they first trained an explanation generator to predict explanations based on the question and answer choices, and then expanded the original classifier input by combining questions and generated explanations. This strategy achieved a significant improvement over the baseline that only used the original data as input. (Kumar and Talukdar 2020) attempted a similar approach, NILE, for the NLI task on the e-SNLI dataset but with a modification where, instead of generating one explanation per training instance, they trained three independent generators conditioned on each label (*entailment*, *neutral*, and *contradiction*), respectively. Then the final NLI model takes as input, the premise-hypothesis pair as well as all three generated explanations. They also evaluated the faithfulness of the explanations to demonstrate that the explanations are well correlated with model predictions, but reported a drop in performance on the NLI task.

Augmenting NLI with Relevant Explanations

The overall workflow of LIREx is shown in Figure 2. Given a premise-hypothesis (P-H) pair, a label-aware rationalizer predicts rationales by taking as input a triplet (P, H, x ; $x \in \{\text{entail, neutral, contradict}\}$) and outputs a rationalized P-H pair, (P, H_x). Next, the NLE generator generates explanations (E_x) for each rationalized P-H pair. Then, the explanations are combined with the original P-H pair as input to the instance selector and inference model to predict the final label. Each component is described below.

Label-aware Rationalizer - $R(\cdot)$

As described in (Camburu et al. 2018), NLEs are created based on the rationales highlighted by annotators. We simulate this process by using a rationalizer to provide the relevant rationales for the NLE generator. This operation models how the explanations are generated for human interpretation.

We formulate this step as a token-level binary classification task where 1 indicates a rationale token and 0 indicates a background token. The rationale classification is only performed on the hypothesis because we consider the premise as background context for the NLI task, where the task is to justify if the hypothesis is an entailment, contradiction, or neutral statement with respect to the premise. Therefore, we hypothesize that the rationales in the hypothesis are sufficient to predict the correct label. We first construct the input sequence as $S^p = \langle s \rangle \text{Label} \langle s \rangle \text{Premise} \langle s \rangle$ and $S^h =$

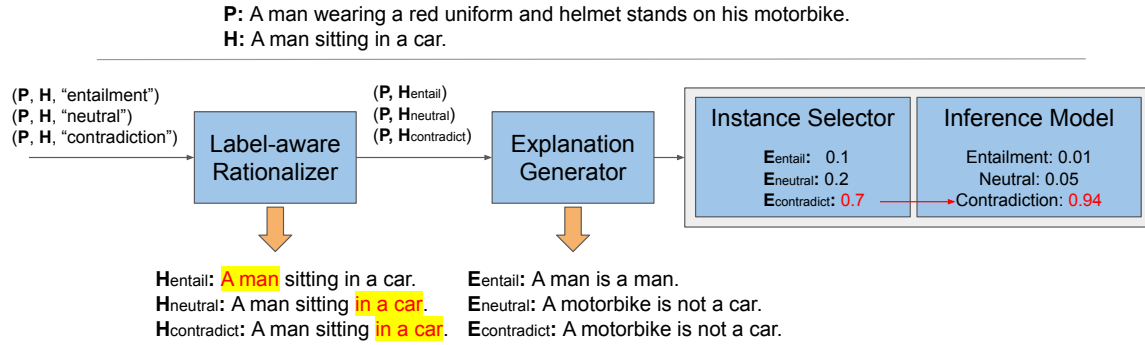


Figure 2: The overall workflow of LIREx framework

$\langle s \rangle$ Hypothesis $\langle s \rangle$, where $\langle s \rangle$ is a special token that separates the components.¹ We append the label information to the premise to inform the rationalizer to highlight label-related rationales. Then we use a RoBERTa_{base} model (Liu et al. 2019) to extract hidden representations for S^p and S^h , denoted as $\mathbf{H}^p = [\dots, \mathbf{h}_i^p, \dots]$ and $\mathbf{H}^h = [\dots, \mathbf{h}_j^h, \dots]$, respectively. Since the rationales in hypothesis depends on the semantic meaning of the premise, we use cross attention to embed premise into hypothesis, defined as

$$a_{ij} = \frac{\exp(\mathbf{h}_i^{hT} \text{Tanh}(\mathbf{W}_1^T \mathbf{h}_j^p))}{\sum_{k=0}^{L^p} \exp(\mathbf{h}_i^{hT} \text{Tanh}(\mathbf{W}_1^T \mathbf{h}_k^p))} \quad (1)$$

$$\hat{\mathbf{h}}_i^h = \text{concat}(\mathbf{h}_i^h, \text{Pool}(\mathbf{H}^p), \sum_k a_{i,k} \mathbf{h}_j^p) \quad (2)$$

where a_{ij} denotes the attention score of the j^{th} token in S^p to the i^{th} token in S^h , L^p denotes the sequence length of S^p , and \mathbf{W}_1 is trainable parameter matrix. Then the new representation of the i^{th} token in S^h is created by concatenating its original state representation, maxpooling representation over \mathbf{H}^p , and the corresponding sum of attentional representation from \mathbf{H}^p . At last, we use a softmax layer with a linear transformation to model the probability of the i^{th} token in S^h being a rationale token: $\text{P}(y_i^h | S^p, S^h) = \text{softmax}(\mathbf{W}_2 \hat{\mathbf{h}}_i^h)$, where \mathbf{W}_2 is a trainable parameter matrix for linear transformation on $\hat{\mathbf{h}}_i^h$.

NLE Generator - G(\cdot)

We model NLE generation as a text generation task, in which we leverage GPT2 (Radford et al. 2019), a language model trained on large-scale language corpus. We choose GPT2_{medium} so that we could have an end-to-end comparison with the previous study that uses the same architecture.

In the previous study (Kumar and Talukdar 2020), the authors finetuned GPT2 independently for each label. Specifically, they trained three GPT2 models separately, which are $G_x(\text{P}, \text{H}, \text{E})$, $x \in \{\text{entail}, \text{neutral}, \text{contradict}\}$. Each G_x is trained only with the P-H pairs annotated as x . As described in the Introduction section, such setup is

¹RoBERTa includes two special tokens, $\langle s \rangle$ and $\langle /s \rangle$. For simplicity, we use $\langle s \rangle$ to denote both of them.

(a) insensitive to the variety of human interpretation toward data, and (b) results in spurious explanations that further harm the label inference task. Additionally, this generation strategy requires training n GPT2 models for n labels, which is still expensive even with fine-tuning. To solve the above issues, we train a single GPT2 model, $G(\text{P}, \text{H}^*, \text{E})$, where H^* is rationalized hypothesis. For example, for the P-H pair in Figure 1b, we construct the input sequence, S^g , as:

```
Premise: P
Hypothesis: A man is [playing] a [musical] [instrument].
Explanation: E
```

where P and E represent the premise and NLE text in training data. To inform the generator about the rationales in hypothesis, we highlight rationale tokens by surrounding them with the square brackets '[']. The generator is finetuned by modeling the text input as a whole. To generate an NLE, we simply remove E from S^g and then use the rest as a text input prompt for the generator.

Unlike the approach in (Kumar and Talukdar 2020) where the label information is appended to the text, we hide the label from the generator to force the model to generate rationale-enabled NLEs. This is consistent with our goal to simulate diverse human interpretation, and prevents the model from generating spurious label-based explanations.

During training and evaluating the explanation generator, we use the rationale tokens and NLEs provided by human annotators. After the generator is trained, we generate three new NLEs for each instance based on the rationales by including each label in S^p , independently. Now each P-H pair is provided with three explanations and we remove the original gold explanations in training data. This is to prevent the instance selector model in the next step from overfitting on the training examples.

Instance Selector and Inference - S(\cdot) and Infer(\cdot)

When a P-H pair and the generated explanations are fed into an inference model, the model benefits from the addition of the explanations when they are correct (cf. Figure 1b). On the other hand, incorrect explanations lead to large uncertainty during the inference process. So, we first select a single plausible explanation for the final inference. To achieve this, we develop a simple strategy assuming that when the

labels are correct, the NLEs generated based on the corresponding enabled rationales are the correct explanations. This allows us to only estimate which NLE is generated by the gold label-enabled rationale. To further simplify this task, if we assume that the gold label-enabled explanation is more likely to be plausible than the other two explanations, we could identify the gold label-enabled explanation by accurately predicting the correct label for the standard NLI task. In other words, a good prediction on the NLE selection task can be achieved by just training a standard NLI classification model.

Training instance selector model We initialize the selector $\mathbf{S}(\cdot)$ with a RoBERTa_{base} model and use the representation of the first token, \mathbf{h}_0 , as the sequence representation. On top of this, an output layer of linear transformation and activation, $\text{Tanh}(\mathbf{U}_1 \mathbf{h}_0) \mathbf{U}_2$, is applied for prediction. \mathbf{U}_1 and \mathbf{U}_2 are parameter matrices. We train $\mathbf{S}(\cdot)$ as a standard supervised learning task where premise and hypothesis are concatenated as a single input sequence, and the model is trained to predict label $Y \in \{\text{entailment, neutral, contradiction}\}$. We pre-train the instance selector and use the label prediction probability distribution as the estimator to find the most likely explanation corresponding to the true label. To improve model robustness, during training, we sample the candidate explanations based on the probability distribution, instead of picking just the most likely explanation. This allows the inference model to better tolerate less plausible explanations. During test phase, we select the explanation with the highest probability.

Training inference model Once the explanation instance is selected, we train the inference model, **Infer**(premise, hypothesis, explanation), with the same model architecture as the selector. Taking insights from the field of weak supervision (Fries et al. 2017; Bach et al. 2017; Ratner et al. 2020) where weakly labeled data is used for training models, we treat the selected explanations as weakly selected instances. Instead of using the standard cross-entropy loss (that requires gold label) as training objective, we use a probability-oriented training objective, soft cross-entropy loss, to improve the model robustness towards noisy input:

$$\text{CE}_{\text{soft}}(\mathbf{p}, \hat{\mathbf{p}}) = \sum_{l \in \{e, n, c\}} \hat{p}_l \log p_l \quad (3)$$

where \mathbf{p} and $\hat{\mathbf{p}}$ are predicted and the “target” probabilities for each label, respectively. In our experiments, we use the estimated probabilities from the instance selector as our “target” probability for training inference model.

Experiments and Results

Data Sets

The proposed framework is evaluated on two widely-used corpora for language inference – SNLI (Bowman et al. 2015) and MultiNLI (Williams, Nangia, and Bowman 2017).

SNLI is a balanced collection of P-H annotated pairs with labels from $\{\text{entailment, neutral, contradiction}\}$. It consists

of about 550K, 10K and 10K examples for train, development, and test set, respectively. (Camburu et al. 2018) recently expanded this data set to e-SNLI in which each data instance is also annotated with explanations. Based on the gold labels for each P-H pair, annotators were asked to highlight the rationale tokens, and provide NLEs based on the rationale. Previous studies (Camburu et al. 2018; Kumar and Talukdar 2020) removed over 17K non-informative training examples (where the explanations contain the entire premise or hypothesis) from their analysis. In our work, to maintain the original training data with minimal changes, we hold out these non-informative training instances only when training the explanation generator, but use the full training data for the remaining steps.

The MultiNLI data set differs from the SNLI data set in that it covers a range of genres of spoken and written text. It contains 433K P-H pairs annotated the same way as SNLI. The evaluation set is divided into Dev-match set (10K) and Dev-mismatch set (10K) – the former is derived from same five domains as the training data and the latter is derived from five other domains.

Model Implementation Details

We compare our results against NILE (Kumar and Talukdar 2020). All NILE models are denoted as NILE_{model} reported by (Kumar and Talukdar 2020) and we designed LIREx_{model} models, LIREx_{model}, accordingly to provide a fair head-to-head comparison.

NILE_{base} and **LIREx_{base}**: The baseline RoBERTa models that takes input only with the P-H pair. **LIREx_{base}** is the reproduced version made for fair comparison.

NILE_{expl} and **LIREx_{expl.*}**: The baseline models that take as input only the explanations. The difference is that LIREx_{expl.*} uses the single explanation selected by the instance selector, while **NILE_{expl}** uses all three explanations.

NILE_{all}: The base model that uses the P-H pair as well as all explanations as input by concatenating them into one input sequence, e.g. “p<s>h<s>e₁<s>e₂<s>e₃”, where all components are separated by a special token <s>.

NILE_{all.extra}: Same as NILE_{all}, except that extra negative samples are created for valid (P, H, E) triplets by sampling from other explanations.

LIREx_{all.max} and **LIREx_{all.prob}**: These two models use the pre-trained instance selector to first select an explanation candidate, and then concatenate the explanation to the P-H pair for input to the inference model. LIREx_{all.max} selects an explanation with the highest probability while LIREx_{all.prob} samples the candidate based on the probability distribution.

Results on In-domain Evaluation

The model performance on SNLI data is summarized in Table 1. We re-implemented the NILE baseline (NILE_{base}) as LIREx_{base} and achieved slight improvement (of 0.06) over published results. The rest of the table shows model improvements compared to the baselines accordingly, with the relative improvements against the baselines summarized in the “vs.baseline” column.

As shown in the table, our final model (LIREx_{all.prob}), is able to achieve an absolute performance gain of 0.32

Model	Dev	Test	vs.baseline	+ Data
SemBERT _{large}	92.0	91.6	-	no
SemBERT _{wwm}	92.2	91.9	-	no
NILE _{base}	91.86	91.49	-	no
NILE _{expl}	88.49	88.11	-3.38	no
NILE _{all}	91.74	91.12	-0.37	no
NILE _{all.extra}	91.29	90.73	-0.76	yes
LIREX _{base}	92.15 \pm .05	91.55 \pm .04	-	no
LIREX _{expl.max}	89.95 \pm .05	89.73 \pm .04	-1.82	no
LIREX _{expl.prob}	90.10 \pm .05	90.03 \pm .05	-1.52	no
LIREX _{all.max}	92.15 \pm .04	91.73 \pm .03	+0.18	no
LIREX _{all.prob}	92.22\pm.03	91.87\pm.03	+0.32*	no

Table 1: Accuracy performance of LIREx on SNLI data (average of five random runs). “*” denotes that the best model is statistically significant (at significance level of 0.05 against baseline and 0.01 against NILE). “+ Data” denotes if additional training data was created. SemBERT (Zhang et al. 2019) is included to refer to the best-reported performance.

Model	Dev-Matched		Dev-MisMatched	
	Acc	vs.baseline	Acc	vs.baseline
NILE _{base}	79.29	-	79.29	-
NILE _{expl}	61.33	-17.96	61.98	-17.31
NILE _{all}	77.07	-2.22	77.22	-2.07
NILE _{all.extra}	72.91	-6.38	73.04	-6.25
LIREX _{base}	80.12	-	79.73	-
LIREX _{expl.max}	65.53	-17.59	65.19	-14.54
LIREX _{expl.prob}	65.57	-17.63	65.32	-14.68
LIREX _{all.max}	79.71	-0.41	79.50	-0.23
LIREX _{all.prob}	79.85	-0.27	79.79	+0.06

Table 2: Transfer performance of LIREx on the out-of-domain MultiNLI data (average of five random runs) without fine-tuning.

accuracy points. In addition, we also provide other variants of our model, namely, LIREX_{expl.max}, LIREX_{expl.prob} and LIREX_{all.max}. All the provided models show that models that use the instance selector achieve better performance. Further, the sampling-based selection strategy (as in LIREX_{*.prob}) performs better than the greedy selection using highest value (as in LIREX_{*.max}).

Results on Out-of-Domain Transfer Evaluation

To test how well our model can generalize to a different data set, we directly apply our model trained on SNLI to an out-of-domain data set, MultiNLI. As shown in Table 2, without any fine-tuning, our model achieves significantly better performance compared to NILE. When compared to the corresponding baselines, our model performance dropped by 0.27 on dev-matched and improved slightly by 0.06 on dev-mismatched. Performance of NILE models, in contrast, dropped significantly on MultiNLI. In addition, compared to the baseline, LIREX_{all.prob}, our final model performs similarly between dev-matched and dev-mismatched, indicating that the inclusion of explanations as supervision improves the overall generalizability of the model.

Model	Dev			Test			Human Eval
	P	R	F1	P	R	F1	
Rationale	59.40	65.22	62.17	59.21	64.89	61.92	90.53

Table 3: F1-based performance of the rationalizer and instance-level accuracy of human evaluation from two annotators over 100 randomly sampled test examples with an inter-rater agreement of 0.89.

Discussion

Behavior of Rationalizer

As described earlier, the explanations are generated based on rationales in the hypothesis. Heuristically, we could evaluate the rationalizer simply by looking at the F1 score to see how well the predictions match the true rationales. However, this evaluation strategy alone is insufficient, because annotators may include additional neighboring tokens as rationales. For example, for the hypothesis “A man sitting in a car” in Figure 1a, “sitting in a car”, “in a car”, “a car”, and “car” are all reasonably correct rationales as they all contain the most important rationale token “car”. If an annotator provides “in a car” as rationale and the rationalizer predicts only “car”, the automated F1 metric will be low even though the main rationale was correctly identified. So, in addition to the F1-based evaluation, we also conducted a manual verification of 100 randomly sampled test examples, and report the instance-level accuracy in Table 3. The manual verification was conducted by two annotators. The annotators were presented with P-H pairs, predicted rationales, and gold rationales, and were asked a “Yes/No” question: Do predicted rationales contain the key information from the gold rationales? Examples annotated as “Yes” were treated as correct predictions. The difference between automated and human evaluation shows that, although the rationalizer did not identify the exact human-provided rationales, it did identify the most important rationales (e.g., “car”) with high accuracy.

An **unexpected, yet preferred** behavior of the rationalizer is that, when there are no obvious rationales towards the pre-appended label information, the model tends to identify rationales that relate to the correct label. For example, for the hypothesis (a) in Table 4, we devised three “alternate” hypotheses to analyze how the rationale predictions change with different hypotheses. We progressively modified a specific component in the original hypothesis – e.g. added “*on stage*” in (b), replaced “*man*” with “*woman*” in (c), and included both these changes and replaced “*musical instrument*” with “*guitar*” in (d). We found that when rationales of a particular label is absent, the rationalizer is prone to output rationales of the correct label (as contradiction rationales in hypothesis (a) and neutral rationales in hypothesis (c)).

Furthermore, when rationales of more than one label exist, the model is capable of identifying at least some correct rationales regarding a label. For example in Hypothesis (b), “playing musical instrument” is an entailment-related rationale, while “playing musical instrument on stage” is a neutral-related rationale. For hypothesis (d), we are able to correctly catch rationales of all labels (“playing” is an entailment, “stage” is neutral and “guitar” is a contradiction).

Premise: A man wearing shirt is playing the drums.		
Hypothesis	Label	Rationales
(a) A man is playing a musical instrument	<u>E</u> N C	musical instrument musical instrument musical instrument
(b) A man is playing a musical instrument on stage	E <u>N</u> C	playing, musical instrument musical instrument, stage playing, musical instrument
(c) A woman is playing a musical instrument	E N <u>C</u>	musical instrument woman, musical instrument woman, musical instrument
(d) A woman is playing guitar on stage	E N <u>C</u>	playing woman, stage woman, guitar

Table 4: Examples of the rationalizer behavior. All hypotheses share the given premise. Rationales are predicted by appending a corresponding label information to the premise. The true labels of the P-H pairs are highlighted with underlines.

Analysis of Explanation Generator

We conducted detailed analyses to show (a) why label information should be removed from generation prompts and (b) robustness of the generator towards rationale variants.

Removing label information from generation prompts prevents the generator from producing too many spurious explanations. In Table 5, we present an example to compare the explanations generated when label information is either included or excluded. For the provided P-H, all three models use GPT2 to generate explanations for each label. **NILE** uses the plain premise and hypothesis combined with label information as generation prompts. **LIREx-w-label** uses the same information as **NILE**, and also includes an indication of rationale tokens in the format described in the NLE Generator section. Finally, in **LIREx-w/o-label**, the label information is removed. As shown in the example, all three models are able to generate reasonable explanations. However, when label information is appended to the input data, spurious explanations are presented (e.g. for E and N labels in **NILE**, and the E label in **LIREx-w-label**). This is because the addition of label serves as a very strong constraint for the generator, which then leads to an explanation that is constructed using the label information as guidance. Even with rationales are provided, the issue could only be slightly alleviated (eg. for N label in **LIREx-w-label**).

However, the issue is resolved when labels are excluded from the generator because the generator is not forced to generate a label-aligned explanation, which could be spurious. Instead, it utilizes only the rationales and the original context to construct the explanations. As discussed earlier, when the rationale of a specified label does not exist, the rationalizer tends to provide rationales related to the true label. So we obtain three reasonable explanations for the E, N, and C labels in **LIREx-w/o-label**(E,N,C). Further, this would greatly increase the chance of a plausible explanation being selected as an explanation candidate by the instance selector.

Premise: Families waiting in line at an amusement park for their turn to ride.			
Hypothesis: People are waiting to see a movie.			
L	NILE	LIREx-w-label	LIREx-w/o-label
Rationale: people, waiting, see, movie			
E	People are waiting to see a movie at an amusement park.	Families are people. If people are waiting in line, it means that they are waiting to see a movie.	People waiting in line are not waiting to see a movie.
Rationale: movie			
N	Not all rides at amusement parks are a movie.	You cannot infer it's for a movie.	Amusement park and movie are not the same.
Rationale: waiting, see, movie			
<u>C</u>	People cannot be waiting to ride an amusement park and a movie at the same time.	People are either waiting to ride an amusement park, or waiting to see a movie.	People are either waiting in line for a ride or waiting to see a movie, but not both.

Table 5: Example of NLEs generated when label information is included and excluded. For each label, we present the generated explanation from each system. “C” (contradiction) is the correct label for the P-H pair.

Robustness of generator towards varied rationales The example in Table 6 shows different explanations generated by using one hypothesis token at-a-time as rationale. Out of the seven generated explanations, only two are invalid. In reality, however, it is rare to have only non-informative stop-words as rationales. This observation strengthens our argument that including rationales in the training of the explanation generator could improve the robustness of the model.

Premise: Families waiting in line at an amusement park for their turn to ride.		
Hypothesis: People are waiting to see a movie.		
Rationale	Explanation	Valid
People are waiting to see a movie	Families implies more than one person are waiting and waiting are different	✓
	You can not infer they are waiting	×
	Rides does not imply to see a movie	✓
	To ride and to see are different	✓
	One cannot ride and see simultaneously.	✓
	Just because families are waiting in line at amusement park doesn't mean they are waiting to see a movie	✓

Table 6: Example of explanations generated using different rationales – one hypothesis token at-a-time as rationale.

Effect of Spurious Explanation

As presented in Tables 5 and 6, **LIREx** is able to consistently generate plausible NLEs with rationale-enabled explanations while **NILE** tends to generate spurious NLEs due

	SNLI-Dev		SNLI-Test		MNLI-M		MNLI-Mis	
	best	rand	best	rand	best	rand	best	rand
LIREx	92.2	91.8	91.8	91.6	79.5	79.5	79.5	79.3
LIREx_N	91.6	86.0	91.5	85.6	79.5	72.0	79.5	72.0

Table 7: Effects of spurious explanations on model performance. LIREx_N uses the same LIREx architecture but with the explanations generated from NILE.

to the inclusion of label information in explanation generation. To show how the spurious explanations could affect model performance, we train our model with the best NLE from the selector for each data instance, and then use a randomly selected NLE during evaluation. The results are presented in Table 7. We observe that since **LIREx** is trained with rationale-enabled NLEs, it suffers only a small performance drop when presented with a randomly selected NLE. On the other hand, if we randomly select an NLE generated by **NILE**, the performance drops significantly compared to when choosing just the best NLE. This shows that (a) **NILE** has a tendency to generate more spurious explanations, and (b) if a spurious explanation is used for training the model, the performance drops significantly. On the other hand, **LIREx** does not use labels when training the generator, and hence, produces fewer spurious explanations, so even a randomly-selected explanation is still relevant.

Faithfulness Evaluation

It is argued by (DeYoung et al. 2019) that a rationale-augmented classifier may not necessarily rely on the rationales but on the original data. Therefore, they propose to measure the faithfulness of the rationales by measuring the *comprehensiveness* (removing rationales from input) and *sufficiency* (using only the rationales as input). Since the **LIREx** inference model uses the generated explanations instead of rationales as input, following (Kumar and Talukdar 2020), we probe the model by removing explanations and using just the explanations to measure comprehensiveness and sufficiency. As shown in Table 8, when compared to the complete input, removal of explanation from the input reduces the performance on all data sets. Just using explanations leads to a significantly larger drop in performance, which is expected because an explanation is more meaningful when combined with the appropriate context (P-H pair) rather than by itself. These two observations show that the model depends on both the P-H pairs and explanations to make predictions, and that the explanations do demonstrate faithfulness. However, it is not clear why the effect on comprehensiveness is not as significant as that on sufficiency.

Relevance Evaluation

Finally, we postulate that trustworthy explanations should be consistent with the appropriate rationale used to interpret the label. Given a specific example that contains different rationales leading a same label, the generator should be able to generate different yet reasonable explanations for each kind of rationale (cf. Figure 1a). We analyze the generated NLEs

	SNLI-Dev	SNLI-Test	MNLI-M	MNLI-Mis
D+E	92.22	91.87	79.85	79.79
D	90.95	91.07	77.10	76.88
E	62.40	62.21	43.35	43.93

Table 8: Faithfulness analysis of LIREx on both SNLI and MultiNLI data. “D+E” uses both P-H pairs and selected explanations as inputs for inference model, “D” uses only the P-H pair, and “E” uses only the selected explanations.

	SNLI-Dev	SNLI-Test	MNLI-M	MNLI-Mis
NILE	84	84	-	-
LIREx	99	97	95	95

Table 9: Manual evaluation of the relevance score over 100 randomly sampled data from each data set by two annotators with the inter-rater agreement of 0.95. NILE evaluations of MultiNLI corpus are missing because we do not have ground truth rationales from human annotators.

based on their **relevance** to human interpretation. From each data set, we randomly sampled 100 examples and ask two annotators to judge the relevance of the generated explanations. Each annotator was provided with context information (premise, hypothesis, rationale, and explanation), and asked to label them as 1 if they agree that the information about the rationales is contained in the explanation, or 0 otherwise. Since **NILE** does not use rationale to generate explanations, we use human-provided rationales in the dataset as the reference target. For **LIREx**, we used predicted rationales as reference targets. As shown in Table 9, **LIREx** is able to maintain a high relevance score between explanations and predicted rationales, even when transferred to the out-of-domain data sets. This shows that the rationale-enabled explanations in **LIREx** are more aligned with human interpretation of the rationales.

Conclusion

In this work, we identified two flaws in the current strategy of using NLEs for the NLI task. To overcome these limitations, we proposed a novel framework, called LIREx, that incorporates both a rationale-enabled explanation generator and an instance selector to augment NLI models with only relevant, plausible NLEs. The code is available at <https://github.com/zhaoxy92/LIREx>. The proposed framework achieves a significant improvement over a strong baseline by 0.32 accuracy points on the SNLI data set, and is comparable to the current state-of-the-art performance on the task. When evaluated over an out-of-domain MultiNLI data set, the proposed approach demonstrated significantly better performance than previously published results without fine-tuning. We conducted extensive qualitative analysis to evaluate each component of our model. Qualitative analysis showed that LIREx generates flexible, faithful, and relevant NLEs that allow the model to be more robust to spurious explanations and better aligned to human interpretation. This work demonstrates the importance and usefulness of including human interpretation in NLI models.

References

- Bach, S. H.; He, B.; Ratner, A.; and Ré, C. 2017. Learning the structure of generative models without labeled data. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, 273–282. JMLR. org.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. *arXiv preprint arXiv:1508.05326* .
- Camburu, O.-M.; Rocktäschel, T.; Lukasiewicz, T.; and Blunsom, P. 2018. e-snli: Natural language inference with natural language explanations. In *Advances in Neural Information Processing Systems*, 9539–9549.
- Camburu, O.-M.; Shillingford, B.; Minervini, P.; Lukasiewicz, T.; and Blunsom, P. 2020. Make Up Your Mind! Adversarial Generation of Inconsistent Natural Language Explanations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 4157–4165.
- DeYoung, J.; Jain, S.; Rajani, N. F.; Lehman, E.; Xiong, C.; Socher, R.; and Wallace, B. C. 2019. Eraser: A benchmark to evaluate rationalized nlp models. *arXiv preprint arXiv:1911.03429* .
- Fidler, S.; et al. 2017. Teaching machines to describe images with natural language feedback. In *Advances in Neural Information Processing Systems*, 5068–5078.
- Fries, J.; Wu, S.; Ratner, A.; and Ré, C. 2017. Swellshark: A generative model for biomedical named entity recognition without labeled data. *arXiv preprint arXiv:1704.06360* .
- Hancock, B.; Bringmann, M.; Varma, P.; Liang, P.; Wang, S.; and Ré, C. 2018. Training classifiers with natural language explanations. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2018, 1884. NIH Public Access.
- Huk Park, D.; Anne Hendricks, L.; Akata, Z.; Rohrbach, A.; Schiele, B.; Darrell, T.; and Rohrbach, M. 2018. Multimodal explanations: Justifying decisions and pointing to the evidence. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8779–8788.
- Kim, J.; Rohrbach, A.; Darrell, T.; Canny, J.; and Akata, Z. 2018. Textual explanations for self-driving vehicles. In *Proceedings of the European conference on computer vision (ECCV)*, 563–578.
- Kumar, S.; and Talukdar, P. 2020. NILE: Natural Language Inference with Faithful Natural Language Explanations. *arXiv preprint arXiv:2005.12116* .
- Lee, D.-H.; Khanna, R.; Lin, B. Y.; Chen, J.; Lee, S.; Ye, Q.; Boschee, E.; Neves, L.; and Ren, X. 2020. LEAN-LIFE: A Label-Efficient Annotation Framework Towards Learning from Explanation. *arXiv preprint arXiv:2004.07499* .
- Li, J.; Miller, A. H.; Chopra, S.; Ranzato, M.; and Weston, J. 2016. Learning through dialogue interactions by asking questions. *arXiv preprint arXiv:1612.04936* .
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* .
- Qin, Y.; Wang, Z.; Zhou, W.; Yan, J.; Ye, Q.; Ren, X.; Neves, L.; and Liu, Z. 2020. Learning from explanations with neural module execution tree. In *International Conference on Learning Representations*.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. *OpenAI Blog* 1(8): 9.
- Rajani, N. F.; McCann, B.; Xiong, C.; and Socher, R. 2019. Explain yourself! leveraging language models for common-sense reasoning. *arXiv preprint arXiv:1906.02361* .
- Ratner, A.; Bach, S. H.; Ehrenberg, H.; Fries, J.; Wu, S.; and Ré, C. 2020. Snorkel: Rapid training data creation with weak supervision. *The VLDB Journal* 29(2): 709–730.
- Srivastava, S.; Labutov, I.; and Mitchell, T. 2017. Joint concept learning and semantic parsing from natural language explanations. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, 1527–1536.
- Wang, S. I.; Ginn, S.; Liang, P.; and Manning, C. D. 2017. Naturalizing a programming language via interactive learning. *arXiv preprint arXiv:1704.06956* .
- Williams, A.; Nangia, N.; and Bowman, S. R. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426* .
- Zhang, Z.; Wu, Y.; Zhao, H.; Li, Z.; Zhang, S.; Zhou, X.; and Zhou, X. 2019. Semantics-aware bert for language understanding. *arXiv preprint arXiv:1909.02209* .
- Zhou, X.; Zhang, Y.; Cui, L.; and Huang, D. 2020. Evaluating Commonsense in Pre-Trained Language Models. In *AAAI*, 9733–9740.