
Knowledge-Grounded Self-Rationalization via Extractive and Natural Language Explanations

Bodhisattwa Prasad Majumder¹ Oana-Maria Camburu² Thomas Lukasiewicz^{3,2} Julian McAuley¹

Abstract

Models that generate extractive rationales (i.e., subsets of features) or natural language explanations (NLEs) for their predictions are important for explainable AI. While an extractive rationale provides a quick view of the features most responsible for a prediction, an NLE allows for a comprehensive description of the decision-making process behind a prediction. However, current models that generate the best extractive rationales or NLEs often fall behind the state-of-the-art (SOTA) in terms of task performance. In this work, we bridge this gap by introducing REXC, a self-rationalizing framework that grounds its predictions and two complementary types of explanations (NLEs and extractive rationales) in background knowledge. Our framework improves over previous methods by: (i) reaching SOTA task performance while also providing explanations, (ii) providing two types of explanations, while existing models usually provide only one type, and (iii) beating by a large margin the previous SOTA in terms of quality of both types of explanations. Furthermore, a perturbation analysis in REXC shows a high degree of association between explanations and predictions, a necessary property of faithful explanations.

1. Introduction

Two approaches that currently predominate for building self-explainable neural models are (i) selecting a subset of input features responsible for a prediction, known as an *extractive rationale* (ER) (Zaidan & Eisner, 2008; Bast-

ings et al., 2019; Sha et al., 2021), and (ii) generating a *natural language explanation* (NLE) for a prediction (Park et al., 2018; Hendricks et al., 2016; Camburu et al., 2018; Kayser et al., 2021). For an explanation (ER or NLE), one is interested in two characteristics: *quality* (or *plausibility*) and *faithfulness*. Quality measures the degree of matching between the model’s explanations and some ground truth; models with low-quality explanations would be undeployable. *Faithfulness* measures how well the explanations reflect the decision-making processes behind the predictions; unfaithful explanations would be misleading.

ERs are concise and provide quick explanations, which may sometimes be enough for users to assess the trustworthiness of the model. However, ERs may not have the means to provide important details of the reasoning of a model (e.g., relations between features) (Wiegrefe et al., 2021). In such cases, NLEs can be complementary, as they allow for detailed justification in a form that is most accessible to humans (natural language). However, machine-generated NLEs, like other generated text, are prone to lacking background knowledge (e.g., commonsense) (Camburu et al., 2020; Mao et al., 2019). This could be because the NLEs are unfaithful or the model did not use the necessary knowledge in its decision-making process. Despite the complementary nature of ERs and NLEs, self-rationalizing models usually provide only one of them, with a few exceptions (Park et al., 2018; Wu & Mooney, 2019). Moreover, while knowledge grounding has been done for black-box models (Bauer et al., 2018; Chandu et al., 2021; Chen et al., 2020a), we are not aware of any work on knowledge grounding for self-rationalizing models. Furthermore, existing self-rationalizing models are often outperformed by black-box models at solving the task at hand, leading to an undesirable performance-explainability trade-off.

To ground both decision-making and rationalization in background knowledge, as well as to reap the benefits of both ERs and NLEs, we combine these three ingredients in a unified self-rationalization framework. Our framework, which we call **REXC** (Extractive Rationales, Natural Language Explanations, and (here) Commonsense)¹, performs five

¹Department of Computer Science and Engineering, UC San Diego, USA. ²Department of Computer Science, University of Oxford, UK. ³Institute of Logic and Computation, TU Wien, Austria. Correspondence to: Bodhisattwa Prasad Majumder <bmajumde@eng.ucsd.edu>.

¹Code is available at <https://github.com/majumderb/rexc>

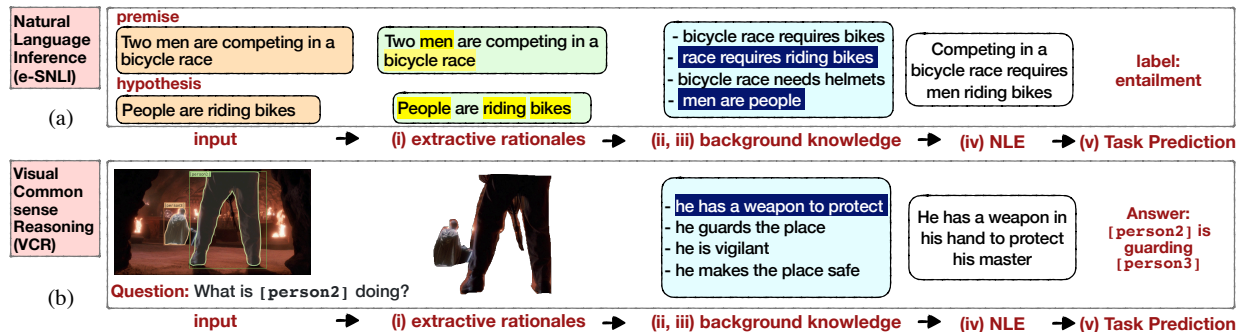


Figure 1. Illustrative examples for REXC on (a) natural language and (b) vision-language tasks.

steps: (i) selects a subset of the input features as an ER, (ii) inputs the ER to a knowledge resource to obtain a set of knowledge snippets about the ER, (iii) selects a subset of the snippets as the most relevant ones for solving the instance, (iv) passes the selected snippets to an NLE generator, (v) passes the generated NLE to a predictor that outputs the final answer (see Figs. 1 and 2). All steps are learned jointly. REXC does not require direct supervision on the ER and snippet selections, which are modeled by two series of latent variables and variational learning (Section 2). Supervision comes from the final answers and NLEs.

REXC is illustrated in Fig. 1. In Fig. 1b, a subset of super-pixels of an input image form the selected ER for the question-answering instance. To answer that “person2 is guarding person3” and explain the answer, the model needs to identify that person2 holds a weapon and have the knowledge that weapons are used to protect.

In our experiments spanning natural language (NL) and vision-language (VL) domains, we find that REXC significantly improves the quality of both ERs and NLEs, while bridging the gap between task performance and explainability. We also show, via perturbation analysis, that the explanations from REXC exhibit necessary conditions of faithfulness. Finally, REXC allows the selection of relevant knowledge snippets even without supervision from the NLEs. As these snippets can act as NLEs, we provide a zero-shot model with NLEs (REXC-ZS), which proves to be competitive with its supervised version.

The contributions of this work are summarized as follows:

- We propose a novel self-rationalizing framework that incorporates background knowledge and provides two complementary types of explanations: ERs and NLEs.
- REXC consistently outperforms previous best models that produce at least one type of explanation and performs on par with the SOTA models that do not provide any explanation, thus bridging the gap between explainability and task performance.
- REXC largely outperforms the previous SOTA in NLE and ER quality.
- REXC passes necessary faithfulness tests.

- REXC allows for a zero-shot setting in terms of NLEs (REXC-ZS), which sometimes outperforms models trained with a full training set of NLEs.

2. REXC

We aim to build a model that solves a task and explains its predictions via both ERs and NLEs. Furthermore, we aim for our model to benefit from resources of background knowledge, which could be general commonsense or domain-specific. To this end, REXC combines these three ingredients in the following way: it extracts rationales from the input, uses them to query an incorporated knowledge module to obtain knowledge snippets, selects the most relevant snippets, generates an NLE, and gives the prediction. We use Fig. 1a as a running example and Fig. 2 for an overview of the architecture.

2.1. Extractive Rationales via Binary Latent Variables

We define a neural module \mathcal{R} that selects an ER from the input. An ER is a minimal sufficient subset of input parts (e.g., tokens for text or super-pixels for images) most responsible for the model’s prediction (Lei et al., 2016). In Fig. 1a, we see an example from the natural language inference task (Bowman et al., 2015) (details in Section 3), where the ER is {“men”, “people”, “bicycle race”, “riding bikes”}, the most responsible units for the prediction (*entailment*).

We model the selection of ERs using a series of latent variables ranging from $[0, 1]$ ($z_i^r \in \mathcal{Z}^r$) over the N input units. A unit becomes a part of the ER *iff* its associated variable takes value 1. Following (Bastings et al., 2019), we use the Hard Kumaraswamy distribution (referred to as HardKuma) as the reparameterization strategy to learn these latent selectors using backpropagation. The parameters of the neural module \mathcal{R} are denoted by θ^r , which estimate the HardKuma variables for the input units. We also encourage the ERs to be terse, and we control the sparsity using an L_1 relaxation defined by the tractable Kumaraswamy CDF.

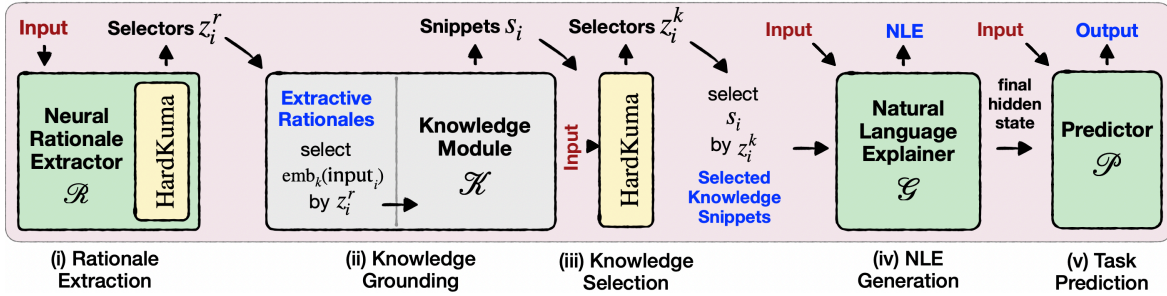


Figure 2. **Architecture of REXC.** The knowledge module is frozen, while the rest of the modules are trained jointly with the signals from the NLEs and outputs. Deliverables from REXC are in blue.

2.2. Knowledge about an Extractive Rationale

We hypothesize that inferred knowledge about the ERs are the most important bits of information for the predictions and, implicitly, for the NLEs. For example, in Fig. 1a, we obtain relevant knowledge snippets (*bicycle race requires bikes and men are people*) for the ER (“bicycle race”, “men”, “people”), which influence both the prediction and the NLE.

We use a knowledge module \mathcal{K} , which supports input from an appropriate modality (e.g., text or image) for querying. We query \mathcal{K} with each contiguous element of the ER (e.g., “bicycle race”) to obtain a large pool of associated knowledge snippets \mathcal{S} . We take advantage of recent developments in generative models capable of providing background knowledge about a given entity for the ease of end-to-end training, such as COMET (Bosselut et al., 2019) for NL inputs and VisualCOMET (Park et al., 2020) for image inputs. The generative knowledge module does not suffer from the no-hit issue that is typically encountered in retrieval settings. However, REXC is flexible to accommodate a retrieval-based knowledge source when equipped with a differential search (see Section 4.4). To facilitate end-to-end training, we use soft representations of the elements of the ER—which are encoded using the embedding layer of \mathcal{K} and subsequently selected by z_i^r (when 1) for queries to \mathcal{K} . Finally, we denote the parameters of \mathcal{K} as θ^k .

2.3. Knowledge Selection

While the knowledge module generates several knowledge snippets (\mathcal{S}), not all of them are relevant for the prediction. Hence, we introduce a knowledge selection step. Furthermore, the selected knowledge snippets can appear as supporting evidence in addition to the generated NLE—an advantage of REXC over models that only generate NLEs.

We model the selection step via another set of latent selectors $z_i^k \in \mathcal{Z}^k$, which take a value from the interval $[0, 1]$ and are realized by a HardKuma distribution (similarly to Section 2.1). More than one knowledge snippet may be relevant, however, we want the knowledge selection to be sparse. Hence, we use L_1 regularization to control the spar-

ity of the selected knowledge. The parameters predicting the latent selectors z_i^k are denoted as θ^{k_s} .

To facilitate end-to-end training, we do not decode knowledge snippets into natural language. Instead, we retain the final hidden representations of each snippet from the knowledge module as $s_i \in \mathcal{S}$. Using z_i^k as an indicator of selection, we obtain the vectors of selected knowledge snippets and concatenate them as input to the NLE generator. We also concatenate the representation of the input for the selector to be able to select the most relevant snippets given the input. At inference time, we decode the selected knowledge snippets into language, which could be used as additional supporting evidence along with the NLE. We call this variant REXC+. Human evaluation shows that this additional evidence leads to higher quality explanations (Section 4.1).

2.4. NLE Generation and Task Prediction

We use a natural language decoder \mathcal{G} , which concatenates the soft representations of the knowledge snippets and of the instance input at the input layer and generates an NLE. After \mathcal{G} , we add a predictor module \mathcal{P} , a linear layer with softmax, which takes the final hidden representation of the NLE and the representation of the instance input, and projects them to the output space for the task prediction. The prediction is thus directly conditioned on the NLE and the input, and, implicitly, on the ER and selected snippets. We denote the parameters of \mathcal{G} and \mathcal{P} as θ^g and θ^p , respectively. We use direct supervision from the ground-truth NLEs and task outputs.

2.5. Training

The parameters for \mathcal{R} , \mathcal{G} , \mathcal{P} , and the knowledge selector can be jointly trained end-to-end with backpropagation by summing up the negative log-likelihoods for the predictions and NLEs. We found that updating parameters for the knowledge resource \mathcal{K} led to a minimal improvement; hence, \mathcal{K} is fixed for computational ease.

However, due to the presence of z_i^r s in \mathcal{R} , we instead

have to optimize a lower bound \mathcal{E} of the original log-likelihood. We follow Bastings et al. (2019) and optimize $\min_{\theta^r, \theta^g, \theta^{ks}, \theta^p} \mathcal{L}_1$ with

$$\mathcal{L}_1 = -\mathcal{E}(\theta^r, \theta^k, \theta^{ks}, \theta^g, \theta^p) + \lambda_0^r \sum_{i=1}^N z_i^r + \lambda_1^r \sum_{i=1}^{N-1} |z_i^r - z_{i+1}^r|, \quad (1)$$

where the second term is the L_1 penalty, the third term is a fused Lasso to control the total number of transitions for compactness (Lei et al., 2016), and λ_0^r and λ_1^r are hyperparameters. Similarly, we have another lower bound for the z_i^k variables in the knowledge selection step, for which we optimize $\min_{\theta^{ks}, \theta^g, \theta^p} \mathcal{L}_2$ with

$$\mathcal{L}_2 = -\mathcal{E}(\theta^{ks}, \theta^g, \theta^p) + \lambda_0^k \sum_{i=1}^M z_i^k, \quad (2)$$

where the second term denotes L_1 regularization for sparse knowledge selection. Finally, we combine the lower bounds as $\alpha \times \mathcal{L}_1 + (1 - \alpha) \times \mathcal{L}_2$, where $\alpha \in [0, 1]$ is a hyperparameter. We estimate the gradient of \mathcal{E} via Monte-Carlo sampling from the reparameterized HardKuma variables (Kingma & Welling, 2014). All hyperparameters are chosen based on a greedy search over the task prediction accuracy (more in Appendix A).

3. Experiments

Tasks. We experiment with three tasks of natural language and two tasks of vision-language understanding as described in Table 1. More task details are in Appendix B.

Table 1. Our tasks: three NL and two VL.

Task	Dataset	Summary
Commonsense Validation	ComVE (Wang et al., 2019)	Choosing input sentence that defies commonsense
Natural Language Inference	e-SNLI (Camburu et al., 2018)	Textual entailment between premise and hypothesis
Commonsense Question Answering	COSe (Rajani et al., 2019)	Answering multi-choice commonsense questions
Visual Entailment	e-SNLI-VE (Kayser et al., 2021)	Entailment between image premise and text hypothesis
Visual Commonsense Reasoning	VCR (Zellers et al., 2019)	Commonsense reasoning in visual question-answering

Implementation Details. The components of REXC for the NL tasks are: **Rationale extraction:** We use the denoising encoder-decoder `bart-base` (Lewis et al., 2020a) with a linear layer and softmax at the end to generate the distribution for latent selectors. **Knowledge source:** We pre-train a `bart-base` model as a proxy for COMET (matched with original perplexity, 11.47 vs. 11.14 as from (Bosselut et al., 2019)) that matches the tokenization scheme used in \mathcal{R} . **NLE and task output:** We use another `bart-base` model to generate the NLEs, decoded with top- p sampling

($p = 0.95$) (Holtzman et al., 2020). A linear layer followed by a softmax is used as the task predictor \mathcal{P} .

The components of REXC for the VL tasks are: **Rationale extraction:** We use a transformer-based VL model, UNITER (Chen et al., 2020b), which uses self-attention to learn contextualized representations for image-text input pairs. We add two MLPs on top of UNITER, which are used to generate the distributions for the latent ER selection from the image and text input; **Knowledge source:** We use VisualCOMET (Park et al., 2020) as an image-based commonsense module, which is fine-tuned on ATOMIC (Sap et al., 2019). For text ERs, we follow the same setup as in the NL setup; **NLE and task output:** We use GPT-2 (Radford et al., 2019), a language decoder, for NLE generation. We adapt GPT-2 to condition on the representations learned by UNITER for VL inputs and use nucleus sampling ($p = 0.95$) for decoding the NLEs. A linear layer followed by a softmax is used for task prediction.

Baselines. We consider existing self-explainable models with the SOTA explanations (NLEs or ERs) as baselines. We also compare REXC with models that are SOTA for task performance (all until now are black-box models for our tasks).

NL Baselines.² The current SOTA for NLEs in all three NL tasks was obtained by WT5 (Narang et al., 2020), a general-purpose NLE generation model. We also compare with works that model NLEs specifically for a dataset: WT5 for ComVE, NILE (Kumar & Talukdar, 2020) for e-SNLI, and CAGE (Rajani et al., 2019) for COSe.

VL Baselines. We compare REXC with: PJ-X (Park et al., 2018) and FME (Wu & Mooney, 2019), two self-rationalizing models that provide both NLEs and ERs, and RVT (Marasovic et al., 2020), a post-hoc explainer that uses external knowledge as REXC. We also compare with e-UG (Kayser et al., 2021), the current SOTA in terms of NLE generation on VL tasks.

Ablations of REXC. We ablate REXC to investigate the effects of each component: ER selector (w/o ER), knowledge selector (w/o KN-Sel), and both (w/o KN & ER). We also ablate with the NLE generator (REXC-ZS), while training just using the final answers as supervision and using the selected knowledge snippets as NLEs. This yields a zero-shot model for NLEs. REXC+ adds the selected knowledge to the NLEs, hence is only used in the human evaluation. Finally, we also investigate the advantage of the generative knowledge module by replacing it with a retrieval-based knowledge source: ConceptNet (Speer et al., 2017) and Visual Commonsense Graphs (Zellers et al., 2019). To make the replacement, we use Maximum Inner Product Search as in (Lewis et al., 2020b). We call this version REXC-RB.

² We used the implementations from the original works.

Table 2. **Task performance (Acc.) and NLE quality for the (a) NL and (b) VL tasks.** NLE Automatic metrics: METEOR (MET.), BERTScore (BRTSc.), BLEURT (BLRT.), and NLE human evaluation metrics: e-ViL score, Yes/No %. **Bold** indicates the best numbers with statistical significance ($p < 0.001$). Underline indicates best task performance from a model with (any type of) explanations.

Model	ComVE						e-SNLI						COSe								
	Acc.	MET.	BRTSc.	BLRT.	e-ViL	Yes No	Acc.	MET.	BRTSc.	BLRT.	e-ViL	Yes No	Acc.	MET.	BRTSc.	BLRT.	e-ViL	Yes No			
Gold	-	-	-	-	91.6	79.3	1.1	-	-	-	-	98.1	94.1	2.7	-	-	-	-	84.8	74.5	1.8
Task SOTA	97.0	-	-	-	-	-	93.1	-	-	-	-	-	-	-	83.7	-	-	-	-	-	-
NLE	-	-	-	-	-	-	-	91.9	11.3	75.3	41.2	84.3	80.1	9.4	-	-	-	-	-	-	-
CAGE	-	-	-	-	-	-	-	-	-	-	-	-	-	-	72.1	1.3	43.1	16.9	59.5	35.4	16.7
WT5	96.1	3.4	86.4	27.0	67.7	46.2	11.0	92.1	12.3	75.3	42.3	85.3	82.7	12.8	81.0	2.2	52.0	22.4	73.0	53.9	10.5
REXC-ZS	96.7	7.7	72.4	24.2	65.8	56.5	16.3	92.4	11.9	63.2	40.7	88.3	85.8	5.5	83.1	2.6	38.1	17.1	83.4	73.2	5.6
REXC	97.2	14.1	91.9	33.7	87.3	72.6	2.8	<u>92.9</u>	19.6	86.8	51.3	94.9	93.9	3.6	83.6	7.2	60.3	30.5	87.4	74.3	2.1
REXC+	<u>97.2</u>	-	-	-	88.4	72.6	1.2	<u>92.9</u>	-	-	-	95.6	94.3	2.7	<u>83.5</u>	-	-	-	87.9	74.7	1.8
REXC-RB	96.4	3.1	89.5	26.1	62.2	43.3	15.1	92.7	13.2	77.4	45.3	87.6	81.2	13.5	82.2	3.7	55.5	23.8	79.3	63.2	9.6
w/o KN-Sel	97.1	11.3	90.2	33.6	84.4	65.3	5.1	92.8	17.9	83.4	51.2	92.8	91.7	5.8	83.2	6.4	58.4	27.9	85.0	70.2	2.5
w/o ER	96.5	5.2	86.1	28.1	67.2	43.4	7.6	92.3	13.1	77.7	43.5	83.4	83.2	15.1	81.4	2.9	52.8	23.8	66.7	45.2	14.9
w/o KN & ER	96.0	4.3	85.2	26.3	66.6	41.3	7.6	92.2	12.4	76.4	41.9	82.9	81.2	15.7	80.8	2.5	51.6	22.4	65.9	44.1	15.9

(a)

Model	e-SNLI-VE						VCR							
	Acc.	MET.	BRTSc.	BLRT.	e-ViL	Yes No	Acc.	MET.	BRTSc.	BLRT.	e-ViL	Yes No		
Gold	-	-	-	-	90.6	79.3	1.1	-	-	-	-	95.8	94.1	2.7
Task SOTA	79.5	-	-	-	-	-	-	81.6	-	-	-	-	-	-
PJ-X	69.2	14.7	79.1	35.6	70.1	55.2	14.5	39.0	16.4	78.4	43.5	73.9	58.2	10.5
FME	73.7	15.6	79.7	34.5	71.9	56.7	13.2	48.9	17.3	79.4	47.8	73.0	56.2	11.1
RVT	72.0	18.8	81.1	35.3	72.2	55.4	12.8	59.0	11.2	78.9	44.2	73.2	57.4	11.5
e-UG	79.5	19.6	81.7	37.8	75.6	57.9	9.9	69.8	11.8	79.0	45.6	75.1	59.3	10.4
REXC-ZS	78.8	12.3	78.6	35.9	79.8	60.7	10.4	79.2	15.8	78.9	41.5	78.9	65.3	10.4
REXC	80.8	22.9	87.7	39.6	81.8	64.2	6.5	<u>79.5</u>	20.9	86.6	53.1	80.9	67.7	7.3
REXC+	80.8	-	-	-	82.1	65.4	6.3	<u>79.5</u>	-	-	-	81.8	67.2	6.2
REXC-RB	78.9	20.7	83.5	38.4	78.3	59.3	10.3	78.9	14.7	81.3	47.2	78.4	62.2	11.4
w/o KN-Sel	79.5	22.4	86.8	39.7	79.9	62.3	7.9	78.6	19.7	85.5	51.4	79.9	67.6	8.2
w/o ER	79.7	20.1	81.9	38.4	76.5	58.6	9.1	74.5	12.4	79.6	46.4	76.3	60.1	10.2
w/o KN & ER	79.4	19.5	81.7	37.7	75.5	57.9	9.8	69.8	11.9	79.0	45.8	75.1	59.4	10.5

(b)

4. Results

4.1. Evaluating the Quality of the Explanations

We evaluate the quality of the ERs and NLEs for REXC in comparison with the baselines.

Automatic Evaluation of NLEs. Following Kayser et al. (2021), we measure the quality of the NLEs by comparing them with the ground truth when the predicted label is correct. Here, we report METEOR (Banerjee & Lavie, 2005), BERTScore (Zhang et al., 2020), and BLEURT (Sellam et al., 2020), which showed the highest correlation with human evaluation (Kayser et al., 2021). More automatic metrics are reported in Appendix C, Table 5.

For NL tasks, REXC achieves the best values on all three automatic metrics (see Table 2a). We see sharp jumps (e.g., ranging from 4.8 to 11 points in METEOR) between REXC and models that do not use knowledge grounding, such as REXC w/o KN & ER and WT5. This confirms that background knowledge is a useful component for better NLEs. The gains for REXC over REXC w/o KN-Sel. show that knowledge selection provides a regularizing effect.

Similarly, REXC outperforms the previous SOTA models

Table 3. **ER quality.** Comparison of previous SOTA models (DeYoung et al., 2020) for rationale extraction vs. REXC for ER quality. Best numbers are in **bold**.

System	e-SNLI			COSe		
	Acc.	IOU	Tok.	Acc.	IOU	Tok.
SOTA	73.4	70.5	70.2	34.6	38.9	51.9
REXC	78.4	72.9	73.5	39.6	41.7	56.1
w/o KN-Sel.	77.8	72.5	73.1	38.7	40.6	55.7

for VL tasks (see Table 2b). In particular, REXC outperforms RVT, a competitive model providing post-hoc NLEs also using the same commonsense resource as REXC, which possibly indicates that joint training for predictions and NLEs is superior over a post-hoc explainability approach.

Automatic Evaluation of ERs. To evaluate the quality of ERs, we directly compare them with gold ERs using ERASER (DeYoung et al., 2020). ERASER uses accuracy (Acc.) and overlap-based metrics such as F1 at Intersection-Over-Union spans (IOU) and token (Tok.) overlap. In Table 3, we show results for e-SNLI and COSe, the only ones from our list that have gold ERs available. We observe that REXC leads to significantly superior-quality ERs compared

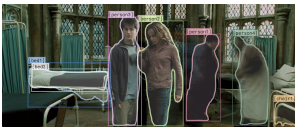

	Input	ER	Knowledge Snippets	NLE	SOTA NLE	Prediction
COSe	<p>Q: People do many things to alleviate boredom. If you can't get out of the house you might decide to do what?</p> <p>A: a) play cards, b) skateboard, c) meet interesting people, d) listen to music</p>	boredom, house, music	<ol style="list-style-type: none"> 1. Music alleviates boredom 2. Music is listened at home 	Music can alleviate boredom when you are alone at home	People listen to music	listen to music
VCR	 <p>Q: Where are [person3] and [person2] right now?</p>	 <p>[person2], [person3]</p>	<ol style="list-style-type: none"> 1. Hospital room has hospital beds 2. Hospital has nurses 	There are hospital beds and nurses in the room	They are patients in the room	They are in a hospital room

Figure 3. Examples of NLEs and ERs generated from REXC along with selected knowledge snippets vs. those from the previous SOTA for the correct predictions for COSe and VCR. Error analysis (Figure 6) and more examples (Figure 9) are in Appendix D.

to models that do not use NLEs or background knowledge to influence rationale extraction (e.g., 56. vs. 51.9 F1). Thus, REXC achieves a new SOTA in ERs for both datasets. Possible explanations for this are: (1) additionally optimizing for NLEs constrains REXC to generate more informative ERs, and (2) to obtain better-suited knowledge snippets, REXC must extract high-quality ERs.

Human Evaluation of NLEs. Following Kayser et al. (2021), we asked human annotators to measure the quality of the generated NLEs. For each NLE, we asked: *Given the input, does the explanation justify the answer?* and provide four options: Yes, Weak-Yes, Weak-No, and No. We report the e-ViL score from (Kayser et al., 2021) combining results for each option with a weight of 1, $\frac{2}{2}$, $\frac{1}{3}$, and 0 respectively. We only evaluate NLEs for correct predictions and collect 250 random such examples for each model and each dataset. More details are in Appendix D.

For NL tasks, Table 2a shows that humans also rated the NLEs from REXC far better than those from the previous SOTA models. Again, REXC without knowledge selection shows large drops, which indicates that the knowledge selection step has positive effects on the quality of the NLEs.

For VL tasks, NLEs from previous SOTA models were rated far lower than ground truths, indicating an even bigger need for improvement. We observe substantial gains for REXC, even when compared to competitive models that already use external knowledge, such as RVT (Marasovic et al., 2020).

Often NLEs generated by REXC are longer than those from the baselines, since they are rich in background knowledge. In the human evaluation sample for e-SNLI, we found that 73% of NLEs from REXC are longer (at least by a token) compared to NLEs from WT5. However, we find that for REXC, length is loosely correlated with the e-ViL score with a Pearson’s correlation score of 0.21. This correlation is similar (0.17) for NLEs from WT5. We also find similarly low correlations (0.13, 0.24, 0.14, and 0.20) between length and e-ViL score for ComVE, COSe, e-SNLI-VE, and VCR, respectively, which indicates that NLE length did not act as

a confounding factor during human evaluation.

Qualitative Analysis. Fig. 3 shows sample outputs from REXC for COSe and VCR (more in Appendix D). We observe that NLEs from REXC are more grounded in knowledge than those from previous SOTA models. Moreover, previous SOTA NLEs fall short of being comprehensive NLEs (e.g., “People listen to music” for COSe), which could be because they do not condition on ERs (e.g., “boredom”).

4.2. Task Performance

Until now, the SOTA models in terms of task performance for all five tasks were models that do not offer any explainability (Wang et al., 2020; 2021; Lan et al., 2020; Xie et al., 2019; Yu et al., 2020). Models that attempt to offer explanations (NLEs or ERs) faced a drop in accuracy (see Tables 2a and 2b). REXC bridges this important gap by matching SOTA task performance for 4 out of 5 tasks and even achieving a new SOTA for e-SNLI-VE, while providing two types of explanations, both of which are of higher quality than the previous models with SOTA explanations.

4.3. Zero-shot NLEs

Often, there exists a high overlap between the generated NLEs and the selected knowledge snippets. This is expected, since the NLEs and predictions are conditioned on the selected knowledge. This raises the question of whether the selected snippets alone could form sufficient NLEs. We argue that, in general, this is not the case, because the information in a background resource may not provide the whole reasoning behind a prediction. This information is only meant to add value but not replace the NLEs. However, in particular cases where the ground-truth NLEs consist mainly of pieces of background knowledge, selected snippets may be sufficient explanations. To investigate this for our datasets, we look at REXC-ZS, where relevant knowledge was selected *only* using the task prediction loss and concatenated to be used as NLEs. Tables 2a and 2b show that REXC-ZS performs poorly in automatic metrics, which

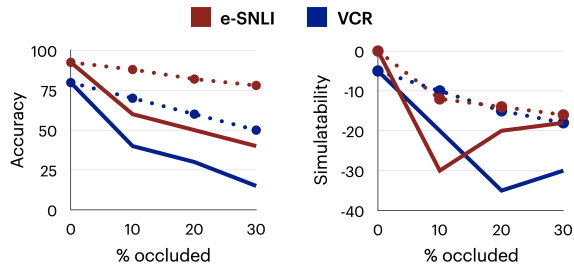


Figure 4. **Feature importance agreement.** Left: Solid lines indicate the prediction accuracy when features important for NLEs are occluded. The dotted lines indicate the prediction accuracy when random features are dropped. Right: solid lines indicate the simulatabilities when features important for prediction are occluded. Dotted lines indicates simulatabilities for random occlusions. In both, solid lines should be lower (meaning higher changes) than dotted lines for better label-NLE association.

is mostly due to being out of distribution w.r.t. the ground-truth explanations. However, in human evaluation, we see that even if the NLEs from REXC-ZS were not better than the generated NLEs from REXC, they were largely better than the NLEs from the previous SOTA models (which were trained with full training sets of NLEs) for 4 out of the 5 tasks. These results indicate that: (1) the NLE module in REXC acts as an important conditional generation step that makes NLEs fluent and more comprehensible; and (2) despite being less fluent, concatenated knowledge snippets can act as NLEs in cases where ground-truth NLEs are not present. This shows the potential of REXC for zero-shot natural language rationalization.

4.4. Generative vs. Retrieval-based Knowledge Module

One of the reasons for choosing a generative knowledge module (COMET and VisualCOMET) is to avoid the no-hit issue of indexed knowledge bases. For example, when we replaced COMET with ConceptNet (Speer et al., 2017), for e-SNLI, we found that 23% of instances do not retrieve any knowledge snippet. As expected, REXC-RB performed worse than REXC (see Tables 2a and 2b).

5. Evaluating Faithfulness

Evaluating the faithfulness of explanations is a challenging open question for both ERs (Jacovi & Goldberg, 2021) and NLEs (Wiegrefe et al., 2021). We analyze REXC for faithfulness based on existing works.

5.1. Faithfulness of the NLEs

Evaluating the faithfulness of NLEs is still in its infancy. To our knowledge, Wiegrefe et al. (2021) is the only work that provides (two) necessary conditions for NLEs’ faithfulness: *feature importance agreement* and *robustness equivalence*. Both conditions perturb the input and measure the change

in model behavior in order to establish the extent of label-NLE association. As they mentioned, there are currently no sufficient conditions for faithful NLEs, since there can be different realizations of NLEs that significantly (but differently) contribute to the model’s prediction process.

Changes in Model Behavior. Change in model behavior can be captured by changes in task accuracy and changes in the predictive ability of NLEs. The predictive ability of NLEs over inputs (formally termed as *simulatability* (Doshi-Velez & Kim, 2017; Hase et al., 2020)) is defined by the change in task accuracy when the generated NLEs are appended to the input. To ensure NLEs’ faithfulness, changes in accuracy and in NLEs (via simulatability) should be similarly affected by changes in the input.

Feature Importance Agreement. This condition uses a gradient-based attribution technique to find the most important features with respect to an output (prediction or NLE). For a predicted class, a gradient attribution is the gradient of the predicted class’s logit with respect to an input feature. The attribution score is calculated by performing an operation (here, L_1 norm) to turn the gradient into a scalar quantity. For REXC, we identify salient input features (tokens or super-pixels) with attribution scores (top- $\{10, 20, 30\}$ %) with respect to the task prediction. We measure the change in simulatability of NLEs when we remove these features from the input. Similarly, we measure the change in task accuracy when we remove the features most important for the NLE generation. To ensure faithfulness, both these changes should be significantly higher than the changes that would appear if we were to remove random input features. Fig. 4 shows that the removal of salient input features similarly affects both task accuracy and NLEs simulatability when compared to random removal—ensuring that this faithfulness condition is met by REXC on e-SNLI and VCR. Similar trends on the other datasets are in Appendix E, Figure 7.

Robustness Equivalence. The second necessary condition involves perturbing the input by adding zero-mean Gaussian noise $\mathcal{N}(0, \sigma^2)$ to the internal representations of its features and observing the corresponding changes in task accuracy and NLE simulatability for a range of noise values. We are interested in noise regions where labels and NLEs remain stable (small changes) and noise regions where labels and NLE become unstable (large changes). To indicate faithfulness of the NLEs, predicted labels and NLEs should remain stable (or unstable) at the same noise region. In Fig. 5, we see this condition holds true for REXC. For example, for e-SNLI (in Fig. 5(a)), we see that the point of minimum contribution of NLEs to the prediction coincides with the sharpest drop in task accuracy, at $\sigma^2 = 25$. Lower noise than $\sigma^2 = 25$ keeps both labels and NLEs stable, whereas higher noise will make both unstable. Similar trends are

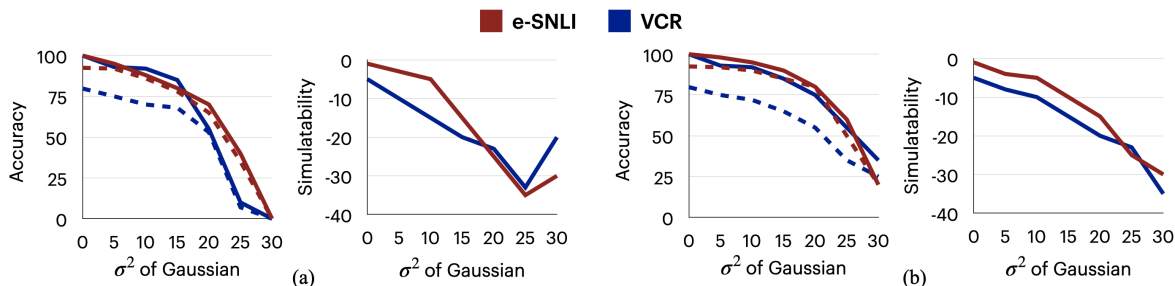


Figure 5. **Robustness equivalence** analysis when noise (with various σ^2) is added to the (a) input and (b) selected knowledge snippets. In each pair, the left chart shows % of stable (unflipped) labels as the solid line, and accuracy of REXC as the dashed line. The right chart in a pair depicts the simulatability of NLEs. For better label-NLE association, the sharpest drop in simulatability and task accuracy should align with the sharpest drop in % of stable labels, so that both the labels and the NLEs are stable (or unstable) in the same noise region.

observed in other datasets (Appendix E, Figure 8).

5.2. Faithfulness of the ERs and Knowledge Snippets

For ERs, faithfulness metrics are more studied than NLEs in the literature (DeYoung et al., 2020; Jacovi & Goldberg, 2021), and both necessary and sufficient conditions for faithfulness exist. DeYoung et al. (2020) introduced two metrics for measuring faithfulness in ERs: *comprehensiveness* (necessary condition) and *sufficiency*. Comprehensiveness is measured by the change in task accuracy between the case when the full input is used for the prediction by the original model and the case when the ERs (from the original model) are dropped (masked for images) and the model is retrained on these new instances (with dropped ERs). A higher difference (maximum 1) would indicate a higher extent of faithfulness. Sufficiency can be calculated as the difference in accuracy between the case when the full input is used for the prediction and the case when *only* the ERs (from original model) are used to retrain the model. A closer to zero value indicates a higher degree of faithfulness. For REXC, we extend this to the selected knowledge snippets to also analyze their comprehensiveness and sufficiency for the task prediction. Table 4 confirms solid comprehensiveness (high values) and sufficiency (close to zero) for both ERs and selected snippets.

A baseline for checking faithfulness of ERs and knowledge selection is to check their sufficiency and comprehensiveness with respect to a *random* selection of input tokens as ER and a random selection of knowledge snippets. Table 4 shows that REXC achieves better comprehensive and sufficiency as compared to a random baseline. REXC also outperforms all models reported in DeYoung et al. (2020) in both metrics.

6. Related Work

Providing explanations for a model’s predictions can be done either post-hoc (via methods that aim to explain already trained and fixed black-box models) or by building

Table 4. **Comprehensiveness (Comp.) and Sufficiency (Suff.)** metrics for ERs and selected knowledge snippets generated by REXC vs. random ERs and knowledge snippets

		ComVE	e-SNLI	COSe	e-SNLI-VE	VCR
ERs						
Random	Comp.	0.12	0.11	0.10	0.13	0.14
REXC	Comp.	0.32	0.45	0.24	0.28	0.33
Random	Suff.	0.44	0.31	0.54	0.51	0.39
REXC	Suff.	0.14	0.08	0.05	0.10	0.13
Knowledge Snippets						
Random	Comp.	0.12	0.14	0.14	0.10	0.09
REXC	Comp.	0.56	0.49	0.36	0.27	0.35
Random	Suff.	0.41	0.51	0.43	0.51	0.37
REXC	Suff.	0.15	0.09	0.08	0.07	0.08

self-explainable models (by jointly producing predictions and explanations). Post-hoc explanations (Lundberg & Lee, 2017; Ribeiro et al., 2016) can be useful when one only has access to a high-performance³ but black-box model. However, post-hoc explanatory methods have been shown to have certain downsides (Adebayo et al., 2018; Slack et al., 2020; Laugel et al., 2019; Camburu et al., 2021; Wiegrefe et al., 2021; Camburu et al., 2019). Moreover, self-explanatory models may benefit from the rich information in the explanations provided at training time (Schramowski et al., 2020; Stacey et al., 2022; Lazaridou et al., 2022). In this work, we focus on self-explainable models to produce two predominant types of explanations: NLEs and ERs.

NLEs. A growing number of works in NL and VL focus on designing neural models that produce NLEs for their predictions to make these models accessible to their users (Hendricks et al., 2016; Camburu et al., 2018; Park et al., 2018; Kayser et al., 2021; Kim et al., 2018; Ling et al., 2017; Marasovic et al., 2020; Wang et al., 2019; Rajani et al., 2019; Zellers et al., 2019). Recently, Narang et al. (2020) achieved SOTA on NLEs for NL tasks by using a pre-trained language model (of 11B parameters, which can be prohibitively large). However, NLEs are sometimes produced separately from

³High performance on held-out sets does not guarantee that the models do the right thing for the right reasons (McCoy et al., 2019).

predictions (Marasovic et al., 2020; Brahman et al., 2021; Atanasova et al., 2020), which raises questions about their faithfulness. In some cases, they were even produced as a task in isolation (without predictions) (Ji et al., 2020). Moreover, the majority of the existing models only produce NLEs, with few exceptions that produce both NLEs and ERs (Park et al., 2018; Wu & Mooney, 2019), as our model does. Furthermore, an analysis on the faithfulness of NLEs is usually missing from the large majority of these works. To our knowledge, only one work recently introduced general necessary conditions for faithfulness in NLEs (Wiegreffe et al., 2021), while few other works attempted architecture-specific faithfulness measures (Kumar & Talukdar, 2020; Wu & Mooney, 2019).

ERs. An early work (Zaidan & Eisner, 2008) investigated rationale extraction from inputs and later was successfully followed by works for both NL (DeYoung et al., 2020; Lei et al., 2016; Bastings et al., 2019; Sha et al., 2021) and VL (Strout et al., 2019) tasks. We model both ERs and NLEs jointly in a novel framework that improves the quality of both types of explanations.

Knowledge Grounding. Free-text generation tasks heavily rely on background knowledge (e.g., commonsense). Several tasks such as dialog generation (Majumder et al., 2020), creative text generation (Chakrabarty et al., 2020; Mao et al., 2019), and counterfactual generation (Bhagavatula et al., 2020) used commonsense for grounding. Recently, Marasovic et al. (2020); Brahman et al. (2021) showed that external knowledge can be useful in separately justifying predictions using NLEs. In this work, we establish that knowledge grounding can be useful in a self-rationalizing framework benefiting both predictions and explanations.

7. Summary and Outlook

In this work, we proposed REXC, a self-rationalizing framework that incorporates background knowledge resources and provides two complementary types of explanations: ERs and NLEs. Using five tasks, from natural language and vision-language domains, we show that REXC obtains a new SOTA performance for both NLEs and ERs. We also close the important gap between task performance and explainability for the five tasks that we experimented with, and obtained a new SOTA for e-SNLI-VE. While we used commonsense resources, future work could look into adding other types of knowledge resources, including more specialized ones, such as legal and medical. Additionally, while we showed that REXC opens up a promising direction for zero-shot NLE generation, further investigation could reap more benefits from the principals behind REXC for zero-shot and few-shot setups.

Acknowledgments

We thank Vered Shwartz, Ana Marasović, the anonymous reviewers and meta-reviewers for their useful comments. Bodhisattwa Prasad Majumder was partly supported by a Qualcomm Innovation Fellowship (2020), UC San Diego Friends of International Center Fellowship (2022), Adobe Research Fellowship (2022), MeetElise, and NSF Award #1750063. Thomas Lukasiewicz and Oana-Maria Camburu were supported by the Alan Turing Institute under the UKRI EPSRC grant EP/N510129/1 and by the UKRI EPSRC grant EP/R013667/1. Thomas Lukasiewicz was additionally supported by the AXA Research Fund and by the ESRC grant “Unlocking the Potential of AI for English Law”.

References

- Adebayo, J., Gilmer, J., Muelly, M., Goodfellow, I., Hardt, M., and Kim, B. Sanity checks for saliency maps. In *NeurIPS*, volume 31, 2018.
- Anderson, P., Fernando, B., Johnson, M., and Gould, S. SPICE: Semantic propositional image caption evaluation. In *ECCV*, pp. 382–398, 2016.
- Atanasova, P., Simonsen, J. G., Lioma, C., and Augenstein, I. Generating fact checking explanations. In *ACL*, pp. 7352–7364, 2020.
- Banerjee, S. and Lavie, A. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization@ACL*, pp. 65–72, 2005.
- Bastings, J., Aziz, W., and Titov, I. Interpretable neural predictions with differentiable binary variables. In *ACL*, pp. 2963–2977, 2019.
- Bauer, L., Wang, Y., and Bansal, M. Commonsense for generative multi-hop question answering tasks. In *EMNLP*, pp. 4220–4230, 2018.
- Bhagavatula, C., Bras, R. L., Malaviya, C., Sakaguchi, K., Holtzman, A., Rashkin, H., Downey, D., Yih, W., and Choi, Y. Abductive commonsense reasoning. In *ICLR*, 2020.
- Bosselut, A., Rashkin, H., Sap, M., Malaviya, C., Celikyilmaz, A., and Choi, Y. COMET: Commonsense transformers for automatic knowledge graph construction. In *ACL*, pp. 4762–4779, 2019.
- Bowman, S. R., Angeli, G., Potts, C., and Manning, C. D. A large annotated corpus for learning natural language inference. In *EMNLP*, pp. 632–642, 2015.

- Brahman, F., Shwartz, V., Rudinger, R., and Choi, Y. Learning to rationalize for nonmonotonic reasoning with distant supervision. In *AAAI*, pp. 12592–12601, 2021.
- Camburu, O., Rocktäschel, T., Lukasiewicz, T., and Blunsom, P. e-SNLI: Natural language inference with natural language explanations. In *NeurIPS*, pp. 9560–9572, 2018.
- Camburu, O., Shillingford, B., Minervini, P., Lukasiewicz, T., and Blunsom, P. Make up your mind! Adversarial generation of inconsistent natural language explanations. In *ACL*, pp. 4157–4165, 2020.
- Camburu, O.-M., Giunchiglia, E., Foerster, J., Lukasiewicz, T., and Blunsom, P. Can I trust the explainer? Verifying post-hoc explanatory methods. In *NeurIPS Workshop Safety and Robustness in Decision Making*, 2019.
- Camburu, O.-M., Giunchiglia, E., Foerster, J., Lukasiewicz, T., and Blunsom, P. The struggles of feature-based explanations: Shapley values vs. minimal sufficient subsets. In *AAAI Workshop on Explainable Agency in Artificial Intelligence*, 2021.
- Chakrabarty, T., Ghosh, D., Muresan, S., and Peng, N. R³: Reverse, retrieve, and rank for sarcasm generation with commonsense knowledge. In *ACL*, pp. 7976–7986, 2020.
- Chandu, K. R., Bisk, Y., and Black, A. W. Grounding ‘grounding’ in NLP. In *Findings of ACL*, pp. 4283–4305, 2021.
- Chen, W., Su, Y., Yan, X., and Wang, W. Y. KGPT: knowledge-grounded pre-training for data-to-text generation. In *EMNLP*, pp. 8635–8648, 2020a.
- Chen, Y., Li, L., Yu, L., Kholy, A. E., Ahmed, F., Gan, Z., Cheng, Y., and Liu, J. UNITER: UNiversal Image-TExt Representation learning. In *ECCV*, pp. 104–120, 2020b.
- Cohen, J. A coefficient of agreement for nominal scales. *Educational and psychological measurement*, 20(1):37–46, 1960.
- DeYoung, J., Jain, S., Rajani, N. F., Lehman, E., Xiong, C., Socher, R., and Wallace, B. C. ERASER: A benchmark to evaluate rationalized NLP models. In *ACL*, pp. 4443–4458, 2020.
- Doshi-Velez, F. and Kim, B. Towards a rigorous science of interpretable machine learning. *arXiv preprint arXiv:1702.08608*, 2017.
- Hase, P., Zhang, S., Xie, H., and Bansal, M. Leakage-adjusted simulatability: Can models generate non-trivial explanations of their behavior in natural language? In *Findings of EMNLP*, pp. 4351–4367, 2020.
- Hendricks, L. A., Akata, Z., Rohrbach, M., Donahue, J., Schiele, B., and Darrell, T. Generating visual explanations. In *ECCV*, pp. 3–19, 2016.
- Holtzman, A., Buys, J., Du, L., Forbes, M., and Choi, Y. The curious case of neural text degeneration. In *ICLR*, 2020.
- Jacovi, A. and Goldberg, Y. Aligning faithful interpretations with their social attribution. *TACL*, pp. 294–310, 2021.
- Ji, H., Ke, P., Huang, S., Wei, F., and Huang, M. Generating commonsense explanation by extracting bridge concepts from reasoning paths. In *AAACL/JCNLP*, pp. 248–257, 2020.
- Kayser, M., Camburu, O., Salewski, L., Emde, C., Do, V., Akata, Z., and Lukasiewicz, T. e-ViL: A dataset and benchmark for natural language explanations in vision-language tasks. In *ICCV*, 2021.
- Kim, J., Rohrbach, A., Darrell, T., Canny, J. F., and Akata, Z. Textual explanations for self-driving vehicles. In *ECCV*, pp. 577–593, 2018.
- Kingma, D. P. and Welling, M. Auto-encoding variational Bayes. In *ICLR*, 2014.
- Kumar, S. and Talukdar, P. P. NILE: Natural language inference with faithful natural language explanations. In *ACL*, pp. 8730–8742, 2020.
- Lan, Z., Chen, M., Goodman, S., Gimpel, K., Sharma, P., and Soricut, R. ALBERT: A lite BERT for self-supervised learning of language representations. In *ICLR*, 2020.
- Laugel, T., Lesot, M.-J., Marsala, C., Renard, X., and Detryniecki, M. The dangers of post-hoc interpretability: Unjustified counterfactual explanations. In *IJCAI-19*, pp. 2801–2807, 7 2019.
- Lazaridou, A., Gribovskaya, E., Stokowiec, W., and Grigorev, N. Internet-augmented language models through few-shot prompting for open-domain question answering. *CoRR*, abs/2203.05115, 2022.
- Lei, T., Barzilay, R., and Jaakkola, T. S. Rationalizing neural predictions. In *EMNLP*, pp. 107–117, 2016.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*, pp. 7871–7880, 2020a.
- Lewis, P. S. H., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., and Kiela, D. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *NeurIPS*, 2020b.

- Lin, C. and Och, F. J. Automatic evaluation of machine translation quality using longest common subsequence and skip-bigram statistics. In *ACL*, pp. 605–612, 2004.
- Ling, W., Yogatama, D., Dyer, C., and Blunsom, P. Program induction by rationale generation: Learning to solve and explain algebraic word problems. In *ACL*, pp. 158–167, 2017.
- Loshchilov, I. and Hutter, F. Fixing weight decay regularization in adam. *CoRR*, abs/1711.05101, 2017.
- Lundberg, S. M. and Lee, S.-I. A unified approach to interpreting model predictions. In *NeurIPS*. 2017.
- Majumder, B. P., Jhamtani, H., Berg-Kirkpatrick, T., and McAuley, J. J. Like hiking? You probably enjoy nature: Persona-grounded dialog with commonsense expansions. In *EMNLP*, pp. 9194–9206, 2020.
- Mao, H. H., Majumder, B. P., McAuley, J. J., and Cottrell, G. W. Improving neural story generation by targeted common sense grounding. In *EMNLP-IJCNLP*, pp. 5987–5992, 2019.
- Marasovic, A., Bhagavatula, C., Park, J. S., Bras, R. L., Smith, N. A., and Choi, Y. Natural language rationales with full-stack visual reasoning: From pixels to semantic frames to commonsense graphs. In *EMNLP Findings*, pp. 2810–2829, 2020.
- McCoy, T., Pavlick, E., and Linzen, T. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *ACL*, 2019.
- Narang, S., Raffel, C., Lee, K., Roberts, A., Fiedel, N., and Malkan, K. WT5?! Training text-to-text models to explain their predictions. *CoRR*, abs/2004.14546, 2020.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. BLEU: A method for automatic evaluation of machine translation. In *ACL*, 2002.
- Park, D. H., Hendricks, L. A., Akata, Z., Rohrbach, A., Schiele, B., Darrell, T., and Rohrbach, M. Multimodal explanations: Justifying decisions and pointing to the evidence. In *CVPR*, pp. 8779–8788, 2018.
- Park, J. S., Bhagavatula, C., Mottaghi, R., Farhadi, A., and Choi, Y. VisualCOMET: Reasoning about the dynamic context of a still image. In *ECCV*, pp. 508–524, 2020.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- Rajani, N. F., McCann, B., Xiong, C., and Socher, R. Explain yourself! Leveraging language models for commonsense reasoning. In *ACL*, pp. 4932–4942, 2019.
- Ribeiro, M. T., Singh, S., and Guestrin, C. “Why should I trust you?”: Explaining the predictions of any classifier. In *SIGKDD*, 2016.
- Sap, M., Bras, R. L., Allaway, E., Bhagavatula, C., Lourie, N., Rashkin, H., Roof, B., Smith, N. A., and Choi, Y. ATOMIC: An atlas of machine commonsense for if-then reasoning. In *AAAI*, pp. 3027–3035, 2019.
- Schramowski, P., Stammer, W., Teso, S., Brugger, A., Shao, X., Luigs, H., Mahlein, A., and Kersting, K. Making deep neural networks right for the right scientific reasons by interacting with their explanations. In *Nature Machine Intelligence*, 2020.
- Sellam, T., Das, D., and Parikh, A. P. BLEURT: Learning robust metrics for text generation. In *ACL*, pp. 7881–7892, 2020.
- Sha, L., Camburu, O., and Lukasiewicz, T. Learning from the best: Rationalizing predictions by adversarial information calibration. In *AAAI*, pp. 13771–13779, 2021.
- Slack, D., Hilgard, S., Jia, E., Singh, S., and Lakkaraju, H. Fooling LIME and SHAP: Adversarial attacks on post hoc explanation methods. In *AIES*, 2020.
- Speer, R., Chin, J., and Havasi, C. ConceptNet 5.5: An open multilingual graph of general knowledge. In *AAAI*, pp. 4444–4451, 2017.
- Stacey, J., Belinkov, Y., and Rei, M. Natural language inference with a human touch: Using human explanations to guide model attention. In *AAAI*, 2022.
- Strout, J., Zhang, Y., and Mooney, R. J. Do human rationales improve machine explanations? *CoRR*, abs/1905.13714, 2019.
- Talmor, A., Herzig, J., Lourie, N., and Berant, J. CommonsenseQA: A question answering challenge targeting commonsense knowledge. In *NAACL-HLT*, pp. 4149–4158, 2019.
- Vedantam, R., Zitnick, C. L., and Parikh, D. CIDEr: Consensus-based image description evaluation. In *CVPR*, pp. 4566–4575, 2015.
- Wang, C., Liang, S., Zhang, Y., Li, X., and Gao, T. Does it make sense? And why? A pilot study for sense making and explanation. In *ACL*, July 2019.
- Wang, C., Liang, S., Jin, Y., Wang, Y., Zhu, X., and Zhang, Y. SemEval-2020 Task 4: Commonsense validation and explanation. In *SemEval*, 2020.
- Wang, S., Fang, H., Khabsa, M., Mao, H., and Ma, H. Entailment as few-shot learner. *CoRR*, abs/2104.14690, 2021.

- Wiegrefe, S., Marasovic, A., and Smith, N. A. Measuring association between labels and free-text rationales. In *EMNLP*, pp. 10266–10284, 2021.
- Wu, J. and Mooney, R. J. Faithful multimodal explanation for visual question answering. In *ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 103–112, 2019.
- Xie, N., Lai, F., Doran, D., and Kadav, A. Visual entailment: A novel task for fine-grained image understanding. *CoRR*, abs/1901.06706, 2019.
- Yu, F., Tang, J., Yin, W., Sun, Y., Tian, H., Wu, H., and Wang, H. ERNIE-ViL: Knowledge enhanced vision-language representations through scene graph. *CoRR*, abs/2006.16934, 2020.
- Zaidan, O. and Eisner, J. Modeling annotators: A generative approach to learning from annotator rationales. In *EMNLP*, pp. 31–40, 2008.
- Zellers, R., Bisk, Y., Farhadi, A., and Choi, Y. From recognition to cognition: Visual commonsense reasoning. In *CVPR*, pp. 6720–6731, 2019.
- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., and Artzi, Y. BERTScore: Evaluating text generation with BERT. In *ICLR*, 2020.

A. Implementation Details

Training. We trained each model for maximum 5 epochs, and training was stopped using an early stopping criteria based on perplexity on the validation sets. For NL tasks, each model is trained with batch size of 4 on two 2080 Ti GPUs. Each REXC variant took 35 hours on ComVE, 45 hours on e-SNLI, and 25 hours on COSe. For VL tasks, each model is trained with batch size of 32 on two 2080 Ti GPUs. Each REXC variant took 85 hours on e-SNLI-VE and 105 hours on VCR.

Hyperparameters. For the rationale extraction step, we set both λ_0^r and λ_1^r to 1.0. This value turned out to be best for both NL and VL tasks. For the knowledge selection step, we set λ_0^g to 0.9, based on validation performance. The α for mixing rationale extraction and NLE generation loss is set to 0.4. We use the AdamW optimizer (Loshchilov & Hutter, 2017) for training each model, and the learning rate was set to $6.25e - 5$, with a linear decay of step size 10^{-1} per epoch. We use BART,⁴ UNITER,⁵ and GPT-2,⁶ with all three being released under the MIT license.

Baselines. We used the official code base for NILE.⁷ For WT5, we fine-tuned a pretrained T5 model.⁸ For all VL baselines (PJ-X, FME, RVT, and e-UG), we followed the implementations details from (Kayser et al., 2021).

B. Tasks

Commonsense Validation. We use ComVE (Wang et al., 2019), a dataset for the task of commonsense validation, where from a pair of sentences, a model needs to choose the sentence that defies commonsense (see Fig. 3). The dataset also comes with NLEs. ComVE consists of 10000/1000/1000 samples in the train/validation/test splits. We use the *BART* tokenizer⁹ to tokenize input strings. The maximum input length was set to 512. The dataset is distributed under the CC BY-SA 4.0 license.

Natural Language Inference. SNLI (Bowman et al., 2015) is a dataset for the task of recognizing textual entailment, where given a pair of sentences (premise and hypothesis), a model must classify their relation as either entailment, contradiction, or neutral. We use the e-SNLI

(Camburu et al., 2018) dataset that contains NLEs for SNLI (see Fig. 3). e-SNLI consists of 550K/10K/10K samples in the train/validation/test splits. We again use the *BART* tokenizer for the input strings. The maximum input length was set to 512. The dataset is distributed under the MIT license.

Commonsense QA. CQA (Talmor et al., 2019) is a multiple-choice commonsense question-answering (QA) dataset. COSe (Rajani et al., 2019) is an extension of CQA that provides an NLE for each correct answer. We treat QA as a multi-class classification task along with generating NLEs for the answer prediction. COSe consists of 9741/1221 samples in the train/validation splits. We use the version 1.11 of the dataset. We use the *BART* tokenizer to tokenize input strings. The maximum input length was set to 1024. The dataset is distributed under the BSD 3-Clause “New” or “Revised” license.

Visual Entailment. SNLI-VE (Xie et al., 2019) is a vision dataset analog to the SNLI dataset (Bowman et al., 2015). SNLI-VE considers an image as a premise (instead of text as in SNLI) and text as a hypothesis, with the same three labels of entailment, neutral, and contradiction. e-SNLI-VE (Kayser et al., 2021) extends SNLI-VE with NLEs. e-SNLI-VE consists of 401K/14K/14K samples in train/validation/test splits. We use the *BERT* tokenization scheme¹⁰ to tokenize text input following UNITER (Chen et al., 2020b). The maximum input length was set to 512. No specific license is associated with the dataset release, and the dataset is freely available.

Visual Commonsense Reasoning. VCR (Zellers et al., 2019) is a dataset for commonsense reasoning in a visual-question-answering setup. We generate the NLEs for each answer prediction from scratch (instead of choosing an NLE from a pool of choices, as the dataset was introduced). VCR consists of 212K/26K/26K samples in train/validation/test splits. Similar to e-SNLI-VE, we use the *BERT* tokenization scheme to tokenize the input text. The maximum input length was set to 512. The license of this dataset is mentioned at <https://visualcommonsense.com/license/>.

C. Automatic Metrics

Following (Kayser et al., 2021), we experiment with a suite of metrics popularly used in language generation to capture how closely the generated NLEs follow the ground truth. We provide additional metrics that were reported in (Kayser et al., 2021), i.e., BLEU-4 (Papineni et al., 2002), ROUGE-L (Lin & Och, 2004), SPICE (Anderson et al., 2016), CIDER (Vedantam et al., 2015) in Table 5.

¹⁰https://huggingface.co/transformers/model_doc/bert.html#berttokenizer

⁴https://huggingface.co/transformers/model_doc/bart.html

⁵<https://github.com/ChenRocks/UNITER>

⁶https://huggingface.co/transformers/model_doc/gpt2.html

⁷<https://github.com/SawanKumar28/nile>

⁸https://huggingface.co/transformers/model_doc/t5.html

⁹https://huggingface.co/transformers/model_doc/bart.html#barttokenizer

Table 5. More Automatic metrics for NL and VL tasks. Best numbers are in **bold** ($p < 0.001$).

	System	BLEU	ROUGE	SPICE	CIDER
ComVE	WT5	21.8	17.2	24.9	34.1
	RExC-ZS	14.5	20.4	16.3	29.4
	RExC-RB	23.6	19.8	27.3	33.5
	RExC w/o KN-Sel	22.1	18.5	24.7	32.3
	RExC w/o KN & ER	21.7	18.2	24.3	31.5
	RExC	25.6	24.5	29.3	37.1
e-SNLI	NILE	29.8	24.3	34.3	47.4
	WT5	32.4	25.3	37.3	48.3
	RExC-ZS	25.5	24.4	33.6	40.1
	RExC-RB	35.6	28.9	39.8	52.5
	RExC w/o KN-Sel	30.5	24.9	35.8	47.3
	RExC w/o KN & ER	31.5	25.3	36.4	48.3
	RExC	37.9	32.4	42.6	54.4
COSe	GPT-2	1.3	8.2	13.2	22.3
	WT5	4.2	12.3	18.3	27.2
	RExC-ZS	3.2	9.4	15.4	23.1
	RExC-RB	3.8	9.8	16.9	27.9
	RExC w/o KN-Sel	4.1	10.9	17.1	26.9
	RExC w/o KN & ER	4.2	11.4	17.3	27.4
	RExC	5.5	18.3	24.3	35.4
e-SNLI-VE	PJ-X	7.3	28.6	24.3	72.5
	FME	8.2	29.9	26.8	83.6
	RVT	9.6	27.3	32.5	81.7
	e-UG	9.6	27.8	34.5	85.9
	RExC-ZS	7.6	24.1	33.2	80.3
	RExC-RB	9.8	26.6	35.1	86.0
	RExC w/o KN-Sel	10.9	27.8	35.9	87.2
	RExC w/o KN & ER	10.1	27.4	35.3	86.1
	RExC	11.2	28.5	36.9	88.2
VCR	PJ-X	3.4	20.5	4.5	19.0
	FME	4.4	22.7	24.2	27.7
	RVT	3.8	21.9	11.7	30.1
	e-UG	4.3	22.5	12.6	32.7
	RExC-ZS	3.6	22.1	25.9	25.6
	RExC-RB	4.9	24.9	28.4	28.2
	RExC w/o KN-Sel	5.3	24.8	28.5	28.3
	RExC w/o KN & ER	5.1	24.4	28.2	27.9
	RExC	5.9	25.4	29.1	29.8

D. Human Evaluation

We designed the human evaluation study based on (Kayser et al., 2021) to assess the NLE quality using Amazon Mechanical Turk. We briefly describe the human evaluation setup here, with a representative snapshot of the UI shown in Fig. 10. For every question, we employed two Anglophone annotators with lifetime HIT acceptance rate of at least 90%.

We made sure that the human annotators are able to solve the predictive task before they evaluate the NLEs. For each NLE, we ask: *Given the input, does the explanation justify the answer?* and provide four options: Yes, Weak-Yes, Weak-No, and No. We report the e-ViL score from (Kayser et al., 2021) combining results for each option. We only consider NLEs for correct predictions and collect 250 random such examples for each model and each dataset. The inter-annotator agreement was captured by Cohen’s Kappa (Cohen, 1960). For each of the datasets, ComVE, e-SNLI,

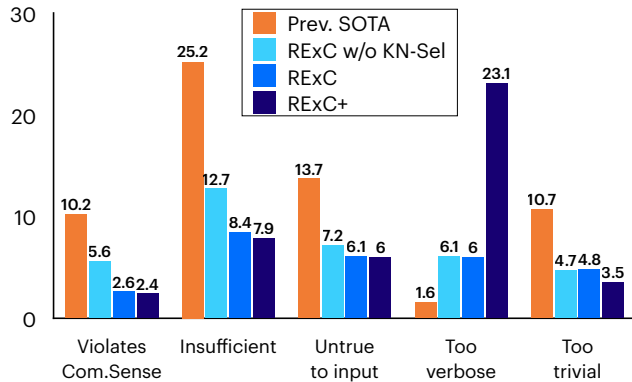


Figure 6. Main limitations of the generated NLEs obtained from user study. All numbers are in % and are averaged by systems and datasets for both NL and VL tasks. Human annotators could choose multiple limitations for an NLE.

COSe, e-SNLI-VE, and VCR, the inter-annotator agreement (kappa) was 0.72, 0.76, 0.79, 0.81, and 0.74, respectively.

Error analysis. Figure 6 summarizes the main drawbacks of generated NLEs (in average) across models and datasets. As main observation, we see that adding commonsense knowledge and knowledge selection in RExC gradually make NLEs more comprehensive and more relevant to the input. While RExC+ wins over all other models across all datasets, human judges often found them too verbose due the presence of supporting knowledge snippets, which might repeat information from the generated NLEs.

Another set of illustrative examples is also given in Fig. 9.

E. Faithfulness

For all datasets, we observe feature importance agreement between labels and NLE, as shown in Fig. 7. Similarly, we see that labels and NLEs are equivalently robust for all datasets, as shown in Fig. 8. This confirms that there exists a strong label-NLE association for RExC—satisfying the necessary conditions for faithful explanations.

Knowledge-Grounded Self-Rationalization via Extractive and NL Explanations

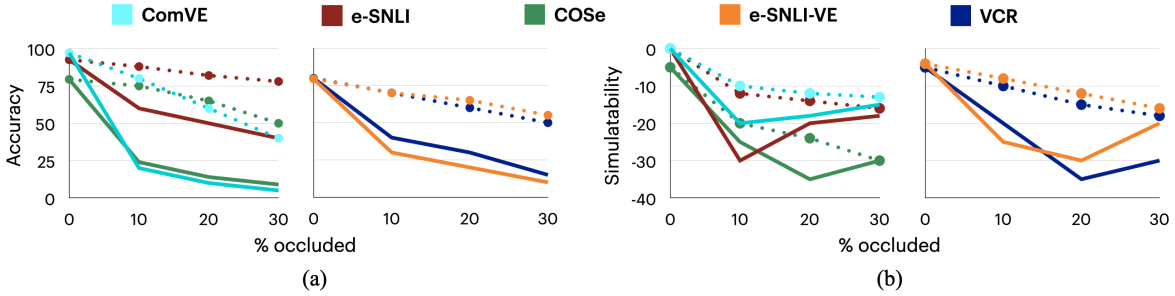


Figure 7. Feature importance agreement with (a) task accuracy and (b) NLE simulatability for all tasks. Details in Section 5.

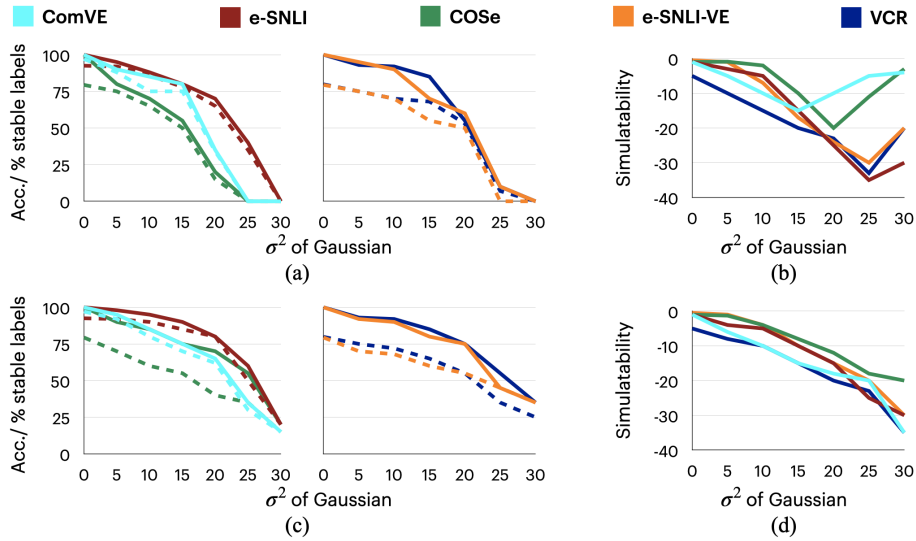


Figure 8. Robustness equivalence analysis when noise (with various σ^2) is added in (a, b) input and (c, d) selected knowledge snippets for all tasks. Details in Section 5.





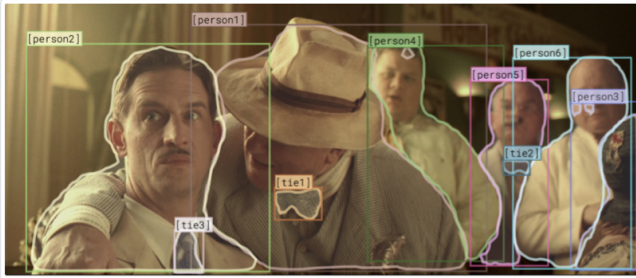
	Input	Rationales	Output	SOTA	REXC	Commonsense ($z_i^* > 0.8$)
ComVE	A: Coffee stimulates people B: Coffee depresses people	coffee	B	Coffee does not depress people	Coffee contains caffeine and is a popular stimulant	1. Coffee contains caffeine 2. Coffee is a stimulant
e-SNLI	Premise: A senior is waiting at the window of a restaurant that serves sandwiches. Hypothesis: A person waits to be served his food.	sandwiches, food	entailment	A person is waiting means a senior is waiting	A person is waiting for sandwiches means a person is waiting for food	1. Sandwich is a food
COSe	Q: Where does a wild bird usually live? A: a) cage, b) sky, c) countryside, d) desert, e) windowsill	wild, bird	sky	Bird flies in the sky	A wild bird flies in free sky	1. Wild bird is free 2. Bird flies in the sky
e-SNLI-VE	 Hypothesis: Some tennis players pose	 tennis, pose	entailment	Some tennis players pose is the same as some tennis players pose.	A group of tennis player are standing together means some tennis players pose	1. A group is standing together 2. Posing is standing pose
VCR	 Q: What is the place?	 place	They are in a hospital room	They are patients in the room	There are hospital beds and nurses in the room	1. Hospital room has hospital beds 2. Hospital has nurses

Figure 9. Examples of NLEs and extractive rationales generated from REXC for all five tasks, along with the pieces of commonsense used by REXC. Generations from the best baseline are included for direct comparison.

Image:



Question: how does [person2] feel about what [person1] is telling him?

1. What is the correct answer?

- He is enjoying it.
- He doesn't like what [person1] is saying.
- He is concerned and a little upset.
- [person6] is upset that [person1] is ridiculing his plan.

Given the image and the question, do the explanations below justify the answer to the question?

Explanation #1: He is in shock thinking something bad is about to happen.

- Yes
- Weak Yes
- Weak No
- No

What are the shortcomings of Explanation #1?

- Contradicts commonsense
- Insufficient justification
- Irrelevant to the inout image and question
- Too verbose or repititive
- Too trivial
- None

It's a good explanation.

Figure 10. Snapshot of our human evaluation with a list of possible shortcomings.