

Towards Interpretable Reasoning over Paragraph Effects in Situation

Mucheng Ren^{1*†} Xiubo Geng^{2†}, Tao Qin³, Heyan Huang¹ and Daxin Jiang^{2‡}

¹School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China

²STCA NLP Group, Microsoft

³Microsoft Research Asia

{renm, hhy63}@bit.edu.cn

{xigeng, taoqin, djjiang}@microsoft.com

Abstract

We focus on the task of reasoning over paragraph effects in situation, which requires a model to understand the cause and effect described in a background paragraph, and apply the knowledge to a novel situation. Existing works ignore the complicated reasoning process and solve it with a one-step “black box” model. Inspired by human cognitive processes, in this paper we propose a sequential approach for this task which explicitly models each step of the reasoning process with neural network modules. In particular, five reasoning modules are designed and learned in an end-to-end manner, which leads to a more interpretable model. Experimental results on the ROPES dataset demonstrate the effectiveness and explainability of our proposed approach.

1 Introduction

As a long-standing fundamental task of natural language processing, machine reading comprehension (MRC) has attracted remarkable attention recently and different MRC datasets have been studied (Rajpurkar et al., 2018; Dua et al., 2019b; Choi et al., 2018; Yang et al., 2018), among which reasoning over paragraph effects in situation (ROPES for short) is a very challenging scenario that needs to understand knowledge from a background paragraph and apply it to answer questions in a novel situation. Table 1 shows an example of the ROPES dataset (Lin et al., 2019), where the background passage states that developmental difficulties could usually be treated by using iodized salt, the situation passage describes two villages using different salt, and questions about which village having more/less people experiencing developmental difficulties need to be answered.

*Work done during internship at STCA NLP Group, Microsoft.

†Equal Contribution

‡Corresponding author

Background

Before iodized salt was developed, some people experienced a number of developmental difficulties, including problems with thyroid gland function and mental retardation. In the 1920s, we learned that these conditions could usually be treated easily with the addition of iodide anion to the diet. One easy way to increase iodide intake was to add the anion to table salt.

Situation

People from two villages ate lots of salt. People from Salt village used regular salt, while people from Sand village people used iodized salt in their diets, after talking to specialists.

Q&A

Q: Which village had more people experience developmental difficulties? A: Salt

Q: Which village had less people experience developmental difficulties? A: Sand

Table 1: An example from the ROPES dataset. Effect property tokens are highlighted in blue, cause property tokens in orange, and world tokens in green.

Almost all existing works (Lin et al., 2019; Khashabi et al., 2020; Dua et al., 2019a; Gardner et al., 2020) for this task adopt a standard MRC approach based on deep learning in one step: the question and a pseudo passage constructed by concatenating the background and situation are fed into a large pre-trained model (e.g. RoBERTa large), and the answer is predicted directly by the model. However, the ROPES task is more complicated than traditional MRC since it requires a model to not only understand the causes and effects described in a background paragraph, but also apply the knowledge to a novel situation. Ignoring the understanding and reasoning process hinders such models from achieving their best performance. Consequently, the best F1 (61.6%) achieved so far

is far below human performance (89.0%). More importantly, such a one-step approach makes the reasoning process unexplainable, which is of great importance for complicated reasoning tasks.

We observe that human solve this kind of complicated reasoning tasks in a sequential manner with multiple steps (Evans, 1984; Sloman, 1996; Sheorey and Mokhtari, 2001; Mokhtari and Reichard, 2002; Mokhtari and Sheorey, 2002). As shown in Table 1, the background paragraph usually states the relationship between a cause property and an effect property, the situation describes multiple worlds each of which is associated with a specific value in terms of the cause property. Human usually does reasoning in a multi-step process: (1) identifying mentioned worlds, (2) identifying the cause and effect property, (3) understanding the relationship between the cause and effect property, (4) comparing identified worlds in terms of the cause property, and (5) reasoning about the comparison of mentioned worlds in terms of the effect property based on (3) and (4).

Inspired by human cognitive processes, in this paper, we propose a sequential approach that leverages neural network modules to implement each step of the above process¹. Specifically, we define

- a *World Detection* module to identify potential worlds,
- an *Effect and Cause Detection* module to identify effect and cause property,
- a *Relation Classification* module to understand the relationship between effect and cause,
- a *Comparison* module to compare identified worlds in terms of the cause property, and
- a *Reasoning* module to infer comparison of mentioned worlds in terms of the effect property.

These modules are trained in an end-to-end manner, and auxiliary loss over intermediate latent decisions further boosts the model accuracy.

Explicitly modeling the sequential reasoning process has two advantages. First, it achieves better performance since the complicated reasoning process is decomposed into more manageable sub-tasks and each module only needs to focus on a

¹The code is publicly available at https://github.com/Borororo/interpretable_ropes.

simple sub-task. Second, intermediate outputs provide a better understanding of the reasoning process, making the learnt model more explainable.

Experimental results on the ROPES dataset demonstrate the effectiveness and explainability of our proposed approach. It surpasses the state-of-the-art model by a large margin (6% absolute difference) in the five-fold cross-validation setting. Furthermore, analyses on intermediate outputs show that each module in our learnt model performs well on its corresponding sub-task and well explains the reasoning process.

2 Related Work

Neural network modules have been studied by several works. Andreas et al. (2016) propose neural module networks with a semantic parser on visual question answering. Jiang and Bansal (2019) apply a self-assembling modular network with only three modules: Find, Relocate and Compare to Hotpot QA (Yang et al., 2018). Gupta et al. (2019) extend the neural module networks to answer compositional questions against a paragraphs of text as context, and perform symbolic reasoning on the self-pruned subset of DROPS (Dua et al., 2019b). Compared with them, we focus on a more challenging MRC task: reasoning over paragraph effects in situation, which has been rarely investigated and needs more complex reasoning. So far as we know, the only two works (i.e. (Lin et al., 2019) and (Khashabi et al., 2020)) on this topic uses a one-step “black box” model. Such an approach performs well on some questions at the expense of limited interpretability. Our work solves this task in a logical manner and exposes intermediate reasoning steps which improves performance and interpretability concurrently.

3 Methodology

As shown in Figure 1, our approach consists of three components which are contextual encoding, interpretable reasoning, and answer prediction.

3.1 Contextual Encoding

We use RoBERTa (Devlin et al., 2019; Liu et al., 2019) to encode background, situation and question together and generate contextualized embeddings. Specifically, given a background passage $B = \{b_i\}_{i=1}^m$, a situation passage $S = \{s_j\}_{j=1}^n$ and a question $Q = \{q_k\}_{k=1}^l$,




Figure 1: The left part is the architecture of our model. The middle part is the interpretable reasoning component in our model. The right part is the summary for inputs and outputs flowing between each module. The encoded contextual representations, H^q, H^s, H^b , serve as global variables for the interpretable reasoning component.

we concatenate them with special tokens as $\langle s \rangle q_1, \dots, q_l \langle /s \rangle \langle /s \rangle s_1, \dots, s_n; b_1, \dots, b_m \langle /s \rangle$, which is then fed into a series of successive transformer blocks contained in RoBERTa,

$$H^q, H^s, H^b = \text{Transformers}(Q, S, B), \quad (1)$$

where $H^b \in \mathbb{R}^{m \times d}$, $H^s \in \mathbb{R}^{n \times d}$, and $H^q \in \mathbb{R}^{l \times d}$ are contextual embeddings for the background, situation, and question, respectively, d is the dimension for hidden states.

3.2 Interpretable Reasoning

World Detection

The module aims to identify concerned worlds from situation according to a question. Take Table 1 as an example, the question cares about two worlds, *Sand Village* and *Salt Village*. To achieve that, we apply a multilayer perceptron (MLP) over the situation representations H^s and normalize the projected logits (using a softmax function) to get attention over all situation tokens for each world,

$$\mathbf{p}_{w_1}^s = \text{softmax}(\text{MLP}(H^s; \theta_{w_1})) \in \mathbb{R}^n, \quad (2)$$

$$\mathbf{p}_{w_2}^s = \text{softmax}(\text{MLP}(H^s; \theta_{w_2})) \in \mathbb{R}^n, \quad (3)$$

where $\mathbf{p}_{w_1}^s$ and $\mathbf{p}_{w_2}^s$ are the attention vectors over situation for the first and second world, θ 's are learnable parameters of MLP. Note that since most examples in the ROPES dataset are related to two concerned worlds, we identify two worlds in our model. However, we can handle multiple worlds by simply extending the module with more MLPs.

Effect and Cause Detection

This module aims to identify effect and cause properties described in the background. To achieve that, another MLP is used to identify the effect property,

$$\mathbf{p}_e^b = \text{softmax}(\text{MLP}(H^b; \theta_e)) \in \mathbb{R}^m. \quad (4)$$

Here \mathbf{p}_e^b is the attention vector over background tokens in terms of the effect property, which attends more to tokens of effect property. Take Table 1 as an example, \mathbf{p}_e^b is the attention over background tokens, whose value is much larger for *developmental difficulties* than other tokens.

Next, we apply a *relocate* operation which re-attends to the background based on the situation and is used to find the cause property in the background (e.g., shifting the attention from *developmental difficulties* to *iodized salt* in Table 1). This is achieved with the help of a situation-aware background-to-background attention matrix $\mathbf{R} \in \mathbb{R}^{m \times m}$,

$$\mathbf{R}_{ij} = \mathbf{w}_{relo}^T \left[(s + \mathbf{H}_i^b); \mathbf{H}_j^b; (s + \mathbf{H}_i^b) \odot \mathbf{H}_j^b \right], \quad (5)$$

$$s = \frac{1}{n} \sum_i^n \mathbf{H}_i^s \in \mathbb{R}^d, \quad (6)$$

where $[\cdot]$ denotes the concatenation operation and \odot is Hadamard product. $\mathbf{w}_{relo} \in \mathbb{R}^{3d}$ is a learnable parameter vector, s can be viewed as an embedding of the whole situation. Then each row of

\mathbf{R} is normalized using the softmax operation. Finally we get the attention vector over background tokens in terms of the cause property \mathbf{p}_c^b ,

$$\mathbf{p}_c^b = \mathbf{R}^T \mathbf{p}_e^b \in \mathbb{R}^m. \quad (7)$$

Here \mathbf{p}_c^b should attend more to the tokens of effect property. For example, *iodized salt* will get larger attention value than other tokens in the background.

Relation Classification

This module aims to predict the qualitative relation between effect and cause property. Take Table 1 as an example, the cause property *iodized salt* and the effect property *developmental difficulty* is negatively correlated. To achieve that, we first derive and concatenate representations of cause and effect property by averaging background representation \mathbf{H}^b weighted by according attention vector, \mathbf{p}_c^b and \mathbf{p}_e^b . Next, we adopt another MLP stacked with softmax to get corresponding probabilities,

$$\mathbf{p}_{rel} = \text{softmax}(\text{MLP}((\mathbf{H}^{bT} \mathbf{p}_e^b; \mathbf{H}^{bT} \mathbf{p}_c^b); \theta_{rel})), \quad (8)$$

where $\mathbf{p}_{rel} = [p_{rel-}, p_{rel+}]$ denotes probability of negative and positive relation, θ_{rel} is a learnable parameter in the MLP. In the example shown in Table 1, p_{rel-} is supposed to be larger than p_{rel+} .

Comparison

This module aims to compare the worlds in terms of the cause property. For example, world 1 (*salt village*) is more relevant to *iodized salt* than world 2 (*sand village*) in Table 1 since people in *salt village* use *iodized salt* while people in *sand village* use *regular salt*.

This is achieved by three steps. First, we derive the attention of cause property over situation \mathbf{p}_c^s from \mathbf{p}_c^b with a similarity matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$ between situation and background,

$$\mathbf{M}_{ij} = \mathbf{H}_i^s \mathbf{W}_{sb} \mathbf{H}_j^{bT}, \quad (9)$$

$$\mathbf{p}_c^s = \mathbf{M} \mathbf{p}_c^b \in \mathbb{R}^n, \quad (10)$$

where $\mathbf{W}_{sb} \in \mathbb{R}^{d \times d}$ are learnable parameters.

Second, we use \mathbf{p}_w^s to mask out irrelevant cause property for each world. This part ensures the alignment between each world and its cause property, which is critical when one situation contains multiple worlds.

$$\mathbf{p}_{c_{w_1}}^s = \text{softmax}(\mathbf{p}_{w_1}^s \odot \mathbf{p}_c^s), \quad (11)$$

$$\mathbf{p}_{c_{w_2}}^s = \text{softmax}(\mathbf{p}_{w_2}^s \odot \mathbf{p}_c^s). \quad (12)$$

Third, each world’s cause property is evaluated by a bilinear function in terms of its relevance to the cause property in background, which is further normalized into a probability with softmax,

$$\text{logit}_{w_1} = (\mathbf{H}^{bT} \mathbf{p}_c^b)^T \mathbf{W}_{com} (\mathbf{H}^{sT} \mathbf{p}_{c_{w_1}}^s), \quad (13)$$

$$\text{logit}_{w_2} = (\mathbf{H}^{bT} \mathbf{p}_c^b)^T \mathbf{W}_{com} (\mathbf{H}^{sT} \mathbf{p}_{c_{w_2}}^s), \quad (14)$$

$$\mathbf{p}_{com_w} = \text{softmax}(\text{logit}_{w_1}, \text{logit}_{w_2}), \quad (15)$$

where $\mathbf{W}_{com} \in \mathbb{R}^{d \times d}$ is a learnable matrix, $\mathbf{H}^{bT} \mathbf{p}_c^b$ represents expected embedding of cause property in background, $\mathbf{H}^{sT} \mathbf{p}_{c_{w_i}}^s$ represents expected embedding of cause property for world i , $p_{com_{w_i}}$ denotes the probability that world i is relevant to cause property.

Reasoning

Given the relationship between effect property and cause property, p_{rel+} and p_{rel-} , and the comparison between worlds in terms of cause property $p_{com_{w_i}}$, this module infers comparison between identified worlds in terms of the effect property. Take Table 1 as an example, given the negative relationship between *developmental difficulties* and *iodized salt*, and *salt village* uses more *iodized salt* than *sand village*, we infer that people in *sand village* are more likely to have *developmental difficulties*.

To this end, we have

$$p_{e_{w_1}} = p_{com_{w_1}} \times p_{rel+} + p_{com_{w_2}} \times p_{rel-}, \quad (16)$$

$$p_{e_{w_2}} = p_{com_{w_1}} \times p_{rel-} + p_{com_{w_2}} \times p_{rel+}, \quad (17)$$

where $p_{e_{w_i}}$ is the probability that world i is more relevant to effect property.

3.3 Answer Prediction

Given intermediate outputs from the interpretable reasoning component, this module predicts the final answer for a question. Specifically, we first convert these intermediate outputs into text spans or 0/1 class as follow.

- We take two steps to convert an attention vector output by *World Detection* or *Effect and Cause Detection* into a text span. First, the token with the highest probability is selected. Then it is expanded with left and right neighbors which are continuous spans and the probability of each token is larger than threshold t . In our experiment we set $t = \frac{1}{l}$, where l is the length of the paragraph.

- For *Comparison*, *Relation Classification*, and *Reasoning*, we select the class with the highest probability.

Then we synthesize a sentence \hat{s} in the format of [World 1] has [larger/smaller] [Effect Property] than [World 2], where choosing “larger” or “smaller” depends on results from the *Reasoning* module. Take Table 1 as an example, the synthetic sentence is *Salt village has larger developmental difficulties than Sand village*. Such synthetic text explicitly expresses comparison between the identified worlds in terms of the effect property. Finally, we concatenate it with the situation s and question q as $\langle s \rangle q; s \langle /s \rangle \langle /s \rangle \hat{s} \langle /s \rangle$, and feed them into RoBERTa which directly predicts the starting and end position of the final answer.

3.4 Model Training

Two models (i.e. interpretable reasoning model; and answer prediction model) are learned in our approach.

Interpretable Reasoning The final loss function for interpretable reasoning is defined as

$$l_{intp} = - \sum_{x \in X} \alpha_x \tilde{x}^T \log(x). \quad (18)$$

Here $X = \{p_{w_1}^s \in \mathbb{R}^n, p_{w_2}^s \in \mathbb{R}^n, p_e^b \in \mathbb{R}^m, p_c^b \in \mathbb{R}^m, p_{c_{w_1}}^s \in \mathbb{R}^n, p_{c_{w_2}}^s \in \mathbb{R}^n, p_{rel} \in \mathbb{R}^2, p_{com_w} \in \mathbb{R}^2, p_{e_w} \in \mathbb{R}^2\}$ are predictions of different modules, $\tilde{x}^T \in \{0, 1\}^n$ or $\tilde{x}^T \in \{0, 1\}^m$ or $\tilde{x}^T \in \{0, 1\}^2$ are corresponding gold labels, and α_x is the weight for module x .

Answer Prediction The training objective of the answer prediction model is defined as

$$l_{ans} = -(\tilde{s} \log(s) + \tilde{e} \log(e)), \quad (19)$$

where $s, e \in \mathbb{R}^{m+n+k}$ are predicted probabilities of the starting and end position, k is the length of the synthetic sentence \hat{s} , and $\tilde{s}, \tilde{e} \in \{0, 1\}^{m+n+k}$ are corresponding gold labels.

4 Experimental Setup

4.1 Dataset

We evaluate our proposed approach on the ROPES (Lin et al., 2019) dataset². So far as we know,

²<https://leaderboard.allenai.org/ropes/submissions/get-started>

Statistics	Train	Dev	Test
background vocabulary size	8,616	2,008	3,988
situation vocabulary size	6,949	1,077	2,736
question vocabulary size	1,457	1,411	1,885
avg. background length	121.6	90.7	123.1
avg. situation length	49.1	63.4	55.6
avg. question length	10.9	12.4	10.6
No. of questions	10,924	1,688	1,710
No. of annotators	7	2	2

Table 2: ROPES statistics

it is the only dataset that requires reasoning over paragraph effects in situation. Given a background paragraph that contains knowledge about relations of causes and effects and a novel situation, questions about applying the knowledge to the novel situation need to be answered. Table 1 shows an example and Table 2 presents the statistics. To be noticed, different from other extractive MRC datasets, train/dev/test set in ROPES is split based on annotators instead of context (Geva et al., 2019; Lin et al., 2019). This might pose a large data bias in each set. For example, as can be seen in Table 2, dev and test sets have similar numbers of questions, while the vocabulary size of background and situation in test set is $2\times$ and $2.7\times$ larger than that in dev set. The same thing happens on the size of question vocabulary, which indicates the existence of the distribution gap between train/dev and test sets and it might lead to underestimate/overestimate the performance of a model.

Cross Validation Because of the limited size of the official dev set and potential data bias between train/dev and test, we conduct 5-fold cross-validation to verify the effectiveness of the proposed approach. K-fold cross-validation assesses the predictive performance of the models and judges how they perform outside the sample to a new data set. Therefore it can assure unbiased results, avoid over-fitting, and testify the generalization capability of a model (Burman, 1989; Browne, 2000; Raschka, 2018). Specifically, we first exclude the labeled 1074 questions from the training data, and then split the remaining training plus dev data into five folds based on background, which ensures each subset has independent vocabulary space and they do not look through each other. For each split, we directly apply modules trained from auxiliary supervision data, and use the training data

to train a model for answer prediction. Averaged results on the 5-fold cross-validation setting are reported.

Auxiliary Supervision We randomly sampled 10% (1,074 questions) of the original training data and labeled them for training the proposed modules. For each example, we label two concerned worlds, effect property in background and situation, values of cause property for two worlds in situation, comparison between two worlds in terms of cause property and effect property, and relationship between cause property and effect property. More detailed guidelines and labeled examples are in Appendix B. Note the neural network modules are trained only on the labeled 1,074 questions.

4.2 Implementation Details

Our model is evaluated based on the pretrained language model RoBERTa large in Pytorch version³. We train the five modules on one P100 16GB GPU and use four GPUs for predicting final answer. We tune the parameter α_x 's according to the averaged performance of all modules, and set it to be 0.05 for span-based loss, 0.2 for the *Comparison* and *Relation* prediction, and 0.3 for the *Reasoning* prediction. Evaluation metrics are EM and F1 which are same as the ones used in SQuAD⁴. The detailed hyperparameters are described in the Appendix A.

4.3 Baseline

We re-implemented the best model (RoBERTa) in the leaderboard⁵ and achieved similar performance (Our implemented baseline achieves EM 69.0 / F1 71.1 on dev and EM 55.2 / F1 61.0 on test while the official one achieves EM 59.7 / F1 70.2 on dev and EM 55.4 / F1 61.1 on test). The basic idea is to leverage the RoBERTa large model (Liu et al., 2019) to encode the concatenation of question, background and situation, and then apply a linear layer to predict the starting and end position of an answer directly.

³<https://github.com/huggingface/transformers>

⁴https://github.com/huggingface/transformers/blob/master/src/transformers/data/metrics/squad_metrics.py

⁵<https://leaderboard.allenai.org/ropes/submissions/public>

Model	Dev		Test	
	EM	F1	EM	F1
RoBERTa* _{Large}	61.4	68.4	64.0	71.1
Ours	73.0	78.1	72.4	77.5
RoBERTa* _{Large} (10% data)	-	-	43.1	53.9
Ours(10% data)	-	-	60.9	71.8
Ours (rule-based)	-	-	54.3	65.5

Table 3: Performance of different models on the ROPES dataset under cross-validation setting.

5 Experimental Results

5.1 Question Answering Performance

Table 3 shows question answering performance of different models, where our approach outperforms the RoBERTa large model by 8.4% and 6.4% in terms of EM and F1 scores respectively. These results show that compared to one-step “black box” model, our interpretable approach which mimics the human reasoning process has a better capability of conducting such complex reasoning.

Furthermore, we also list the performance of our approach and the baseline model when using only randomly sampled 10% of training data in Table 3. That is, both the neural network modules and answer prediction model in our approach are trained with only 1074 questions. As seen in the table, our model learned from 10% of training examples achieves competitive performance to the baseline model learned from full data (71.8% v.s. 71.1% in terms of F1 score). In contrast, the performance of the baseline model drops dramatically by 32%. This indicates that traditional black-box approach requires much more training data while our approach has better generalization ability and can learn the reasoning capability with much fewer examples.

We also implement a rule-based answer prediction approach (detailed descriptions in the Appendix E), which are generated based on the same 10% of training examples as in interpretable reasoning components. As shown in Table 3 the rule-based approach performs worse than the RoBERTa model, indicating better generalization ability of pre-trained models.

5.2 Case Study of Interpretability

The most remarkable difference between our model and the one-step “black box” model is that our

Background	
Storing large volumes of data - When storing XML to either file or database, the volume of data a system produces can often exceed reasonable limits, with a number of detriments: the access times go up as more data is read, CPU load goes up as XML data takes more power to process, and storage costs go up....	
Situation	
Tory had a busy day storing XML. At 7 AM, he stored 201 Gigabytes to the database. At 8 AM, he stored 301 Gigabytes to the database. At 9 AM,... At 10 AM, ... At 11 AM, ... At 12 PM, ... At 1 PM, he went to sleep to finish storing XML later on that day.	
Q&A	
Q: What time did CPU load go up: 8 AM or 1 PM? A: 8 AM	
Predictions	
Worlds:	[8 AM, 1 PM]
Effect:	CPU load goes up
Cause ^B :	Storing large volumes of data
Cause ^S _{World1} :	301 Gigabytes
Cause ^S _{World2} :	sleep
Cause Cmp:	World 1
Relation:	positively related
Effect Cmp:	World 1
Final Answer:	8 AM

Table 4: A running example with visualized intermediate outputs of our approach.

model outputs multiple intermediate predictions which well explains the reasoning process. Note all modules in our model output probabilities or attention on input text, which are further fed into downstream ones for end-to-end learning. In order to explicitly visualize the output of each module, we take a similar approach to §3.3 to convert these probabilities into a text span or a 0/1 classification.

We demonstrate the reasoning process of our model with a running example shown in Table 4. Please see more examples in Appendix D. Here the background states the relationship between CPU load and data volume, i.e. CPU load goes up when processing larger volume of data. The situation describes that Tory stored different sizes of data at a different time. For example, he stored 301 Gigabytes at 8 AM and went to sleep at 1 PM. Finally, the question asks to compare CPU loads between 8 AM and 1 PM. As shown in Table 4, our model

	F1	Fuzzy F1		Accuracy
World 1	83.5	86.8	Comparison	83.8%
World 2	84.4	86.1	Relation	84.5%
Effect	67.8	83.6	Reasoning	74.0%
Cause ^B	57.6	70.1		
Cause ^S _{World1}	69.4	81.3		
Cause ^S _{World2}	58.2	71.9		

Table 5: Performance of Each Module

outputs several intermediate results. First, it identifies two concerned worlds, 8 AM and 1 PM from the situation. Then it predicts the effect property, CPU load goes up, given which the cause property in the background (i.e. storing large volumes of data) and according values for the two worlds (i.e. 301 Gigabytes and sleep) are predicted. Next, it compares the two worlds in terms of cause property and predicts that world 1 is larger than world 2. Also it predicts that the cause property and effect property is positively related, i.e. the relation is classified as 1. Finally, it reasons that world 1 takes higher CPU loads than world 2. This example demonstrates that our approach not only predicts the final answer for the question, but also provides detailed explanations for the reasoning process.

5.3 Neural Network Module Performance

Taking the same approach as in §5.2, we convert the output of each module into a text span or a predicted class. We manually sampled another 5% of the training data, labeled them with outputs for each module, and evaluate the visualized results of all modules. Table 5 summarizes the performance for each module, where the predicted text span is measured by F1 score and classification prediction is measured by accuracy.

World Detection This module implements a similar capability as traditional extractive MRC, since both require to detect concerned text spans from a passage according to a question. Consequently, it achieves similar performance to top models of the popular SQuAD dataset⁶, where our World Detection module reaches about 83% F1 score and single RoBERTa large model on SQuAD gets about 89%. The gap might come from different modeling styles. Our model predicts the probability of each token being concerned, while SQuAD models directly predict the starting and end position of an answer,

⁶<https://rajpurkar.github.io/SQuAD-explorer/>

which performs better on boundary detection.

Effect and Cause Detection Compared with *World Detection*, the F1 score for this module decreases but actually still acceptable (F1=67.6 for Effect^B, 57.6 for Cause^B). The most possible reason is that, effects and causes are usually longer than world names. For example, the average length of world name is 1.2, while those of effect and cause are 2.7 and 2.2 respectively. Longer text span increases the difficulty of prediction.

For above two span-related modules, we argue that since our model leverages the attention score in a soft way, it is less sensitive to accuracy of boundary changes. Therefore, we added another fuzzy F1 score for them. The fuzzy F1 of each question is set to 1 as long as its original F1 is larger than 0. As shown in Table 5, the fuzzy F1 scores of these two modules increase to 70%~86%, indicating good reasoning capability of them.

Comparison, Relation Classification These two modules essentially requires the capability of classification. The high accuracy (83.8% and 84.5%) indicates that our modeling approach can effectively leverages the prediction of upstream modules and does a good job on them.

Reasoning Given the high accuracy of the *Comparison* and *Relation Classification* modules, the *Reasoning* model achieves 74% of accuracy, which provides high-quality input for final answer prediction.

5.4 Error Analysis

We randomly sampled 200 wrongly predicted questions to do error analysis and find that they fall into two major types, which are described below (Please see complete description of questions, backgrounds and situations in the Appendix C).

5.4.1 Type One Error

Type one errors are caused by wrong model predictions, most of which occur in below three modules.

Wrong Predicted Worlds Such errors are mainly caused by length imbalance between question and situation. Since situation is usually much longer than question, the *World Detection* module might make the same predictions for different questions of same situation.

Wrong Predicted Cause Property in Situation Such errors are mainly caused by imbalanced de-

scriptions for different worlds, where a situation describes details for one world but mentions another world with very few words. In such cases, the *Comparison* module might assign the same cause property for different worlds in situation.

Wrong Predicted Comparison Results Such errors often occur when two worlds are described with similar words, e.g. “high” v.s. “higher”, or “smoking” vs. “not smoking”, in which case the *Comparison* module might be confused by similar expressions of two worlds and fail to compare them.

5.4.2 Type Two Error

Type two errors occur when the proposed framework is not suitable to solve the questions. Here we list some example cases.

Missing Knowledge Background paragraph does not provide sufficient knowledge for reasoning. For example, a background paragraph only describes information about fish while the questions asks fertilization take place inside or outside of mother’s body for a mammal creature.

Implicit Worlds Concerned worlds in a question are not explicitly described in the situation. For example, a situation paragraph says that Matthew does intensive workouts while the question asks that his strengths will increase or decrease when he stops working out. In such a case, the *world* that Matthew stop working out is not explicitly described in the situation.

Additional Math Computation Answering such questions requires additional math computation. For example, a background states the speed of sound waves in air/water/iron and the question asks how much faster a channel (with water) would be than another channel (with air). Answering such questions requires additional math computation (i.e. subtraction, addition etc.)

6 Conclusion and Future Work

In this paper, we aim to answer ROPES questions in an interpretable way by leveraging five neural network modules. These modules are trained in an end-to-end manner and each module provides transparent intermediate outputs. Experimental results demonstrate the effectiveness of each module, and analysis on intermediate outputs presents good interpretability for the inference process in contrasted with “black box” models. Moreover, we

find that with explicitly designed compositional modeling of inference process, our approach with a few training examples achieves similar accuracy to strong baselines with full-size training data which indicates a better generalization capability. Meanwhile, extending these models to a larger scope of question types or more complex scenarios is still a challenge, and we will further investigate the trade-off between explainability and scalability.

Acknowledgement

We acknowledge this work is supported by National Natural Science Foundation of China (No.61751201) and National Key RD Plan (No.2016QY03D0602), We would also like to thank the anonymous reviewers for their insightful suggestions.

References

- Jacob Andreas, Marcus Rohrbach, Trevor Darrell, and Dan Klein. 2016. Neural module networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 39–48.
- Michael W Browne. 2000. Cross-validation methods. *Journal of mathematical psychology*, 44(1):108–132.
- Prabir Burman. 1989. A comparative study of ordinary cross-validation, v-fold cross-validation and the repeated learning-testing methods. *Biometrika*, 76(3):503–514.
- Eunsol Choi, He He, Mohit Iyyer, Mark Yatskar, Wentau Yih, Yejin Choi, Percy Liang, and Luke Zettlemoyer. 2018. [QuAC: Question answering in context](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2174–2184, Brussels, Belgium. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dheeru Dua, Ananth Gottumukkala, Alon Talmor, Sameer Singh, and Matt Gardner. 2019a. Orb: An open reading benchmark for comprehensive evaluation of machine reading comprehension. *arXiv preprint arXiv:1912.12598*.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019b. [DROP: A reading comprehension benchmark requiring discrete reasoning over paragraphs](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2368–2378, Minneapolis, Minnesota. Association for Computational Linguistics.
- Jonathan St BT Evans. 1984. Heuristic and analytic processes in reasoning. *British Journal of Psychology*, 75(4):451–468.
- Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, et al. 2020. Evaluating nlp models via contrast sets. *arXiv preprint arXiv:2004.02709*.
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets. *arXiv preprint arXiv:1908.07898*.
- Nitish Gupta, Kevin Lin, Dan Roth, Sameer Singh, and Matt Gardner. 2019. Neural module networks for reasoning over text. *arXiv preprint arXiv:1912.04971*.
- Yichen Jiang and Mohit Bansal. 2019. [Self-assembling modular networks for interpretable multi-hop reasoning](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4474–4484, Hong Kong, China. Association for Computational Linguistics.
- Daniel Khoshdel, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. Unifiedqa: Crossing format boundaries with a single qa system. *arXiv preprint arXiv:2005.00700*.
- Kevin Lin, Oyvind Tafjord, Peter Clark, and Matt Gardner. 2019. [Reasoning over paragraph effects in situations](#). In *Proceedings of the 2nd Workshop on Machine Reading for Question Answering*, pages 58–62, Hong Kong, China. Association for Computational Linguistics.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Kouider Mokhtari and Carla A Reichard. 2002. Assessing students’ metacognitive awareness of reading strategies. *Journal of educational psychology*, 94(2):249.
- Kouider Mokhtari and Ravi Sheorey. 2002. Measuring esl students’ awareness of reading strategies. *Journal of developmental education*, 25(3):2–11.

Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.

Sebastian Raschka. 2018. Model evaluation, model selection, and algorithm selection in machine learning. *arXiv preprint arXiv:1811.12808*.

Ravi Sheorey and Kouider Mokhtari. 2001. Differences in the metacognitive awareness of reading strategies among native and non-native readers. *System*, 29(4):431–449.

Steven A Sloman. 1996. The empirical case for two systems of reasoning. *Psychological bulletin*, 119(1):3.

Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D. Manning. 2018. [HotpotQA: A dataset for diverse, explainable multi-hop question answering](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2369–2380, Brussels, Belgium. Association for Computational Linguistics.

Appendices

A Parameters List

Interpretable Reasoning	Search Space(Bounds)	Best Assignment
No. of GPU(P100)	1	1
Average runtime (mins)	45	45
No. of params.(include LM)	361661706	361661706
No. of Layers in MLP	3	3
No. of Search trials	38	38
learning rate optimizer	Adam	Adam
Max Seq. Length	<i>choice</i> [384,512]	512
Doc stride	<i>choice</i> [64,128]	64
Learning Rate	<i>uniform-float</i> [5e-6,3e-5]	2e-5
Batch Size per GPU	<i>choice</i> [1,2,4]	2
Gradient Accumulation Step	<i>choice</i> [1,2]	1
No. of Epoch	<i>uniform-integer</i> [1,5]	4
Fixed Length for Q,S,B	<i>uniform-integer</i> [20,30],[150,250],[350,450]	30,200,400

Table 6: Detailed parameters used in Interpretable Reasoning, we provide search bounds for each hyperparameter and list out the hyperparameters combination for out best model. Other unmentioned parameters keep same as the one used in BERT.

B Auxiliary Supervision Instruction

Table 8 shows one labelled example, and the process of adding auxiliary supervision label contain the following steps:

1. Annotate the samples manually: For the selected examples, we find the spans for the worlds, cause property and effect property in the background, cause property for the Worlds in the situation and decide the results for comparison, relation and reasoning modules.

Answer Prediction	Search Space(Bounds)	Best Assignment
No. of GPU(P100)	4	4
Average runtime (mins)	60	60
No. of params.(include LM)	355361794	355361794
No. of Search trials	64	64
learning rate optimizer	Adam	Adam
Max Seq. Length	<i>choice</i> [384,512]	384
Doc stride	<i>choice</i> [64,128]	128
Learning Rate	<i>uniform-float</i> [5e-6,3e-5]	1.5e-5
Batch Size per GPU	<i>uniform-integer</i> [4,8]	4
Answer Length Limit	<i>uniform-integer</i> [5,30]	9
Gradient Accumulation Step	<i>choice</i> [1,2]	1
No. of Epoch	<i>uniform-integer</i> [1,5]	4

Table 7: Detailed parameters used in Answer Prediction, we provide search bounds for each hyperparameter and list out the hyperparameters combination for out best model and baseline model. Other unmentioned parameters keep same as the one used in BERT.

Background

As a cell grows, its volume increases more quickly than its surface area. If a cell was to get very large, the small surface area would not allow enough nutrients to enter the cell quickly enough for the cell’s needs...Such cell types are found lining your small intestine, where they absorb nutrients from your food through protrusions called microvilli .

Situation

There are two cells inside a Petri dish in a laboratory, cell X and cell Z. These cells are from the same organism, but are not the same age. Cell X was created two weeks ago, and cell Z was created one month ago. Therefore, cell Z has had two extra weeks of growth compared to cell X.

Q&A

Q: Which cell has a larger volume?

A: cell Z

Labels

World1:	<i>Cell X</i>	[142,148]
World2:	<i>cell Z</i>	[180,186]
Effect:	<i>volume</i>	[21,27]
Cause ^B :	<i>cell grows</i>	[5,15]
Cause ^S _{World1} :	<i>two weeks ago</i>	[161,174]
Cause ^S _{World2} :	<i>one month ago</i>	[199,212]
Cause Cmp:	<i>cell Z</i>	1
Relation:	<i>positively related</i>	1
Effect Cmp:	<i>World 2</i>	1

Table 8: An example with auxiliary supervision labels.

2. Generate machine-readable labels automatically: Then we use scripts to automatically transform the annotations to the machine-

readable form, i.e. we record the start and end character index for all spans and keep the results for comparison, relation and reasoning modules as binary form.

C Error Cases for Modules

Table 9 lists out several type 1 error cases mentioned in the Error Analysis part, while Table 10 lists out type 2 error cases which beyonds the scope of our model.

D More Examples

We present more examples that correctly answered by our model in Table 11.

E Heuristic Rules for Answer Prediction

We also conduct a rule-based approached to predict the final answer which contains the following steps:

1. Categorize questions based on the type of answer: By looking at the labeled train dataset, we can summarize that the types of answer can be divided into two types: 1) World Type, answer is one of the compared worlds; 2) Comparative Word Type, like "more" or "less".
2. For World Type, we filter out such type of questions by searching question keywords, for example, questions started with { *What, Which, Who, Where, When* } usually have world type of answers. Then we determine the results based on the prediction obtained in Reasoning Module.
3. For Comparative Word Type, we further filter out this type of questions from the rest questions by defining a list of comparative word pair like { '*more*': '*less*', '*higher*': '*lower*' ... }. Then we identify the primary world that being compared in the question and associate it with our identified worlds from Group Detection module, then determine the comparative word for the primary compared world by using the results from Reasoning module.
4. For the remaining questions, we simply return the world with higher effect property probability from Reasoning module as the final answer.

Examples	Prediction
<p>ID:1867731649 & 350571026</p> <p>Background: Fish mortality is a parameter used in fisheries population dynamics to account for the loss of fish in a fish stock through death. The mortality can be divided into two types:Natural mortality: the removal of fish from the stock due to causes not associated with fishing. Such causes can include disease, competition, cannibalism, old age, predation, pollution or any other natural factor that causes the death of fish. In fisheries models natural mortality is denoted by (M).[1]Fishing mortality: the removal of fish from the stock due to fishing activities using any fishing gear It is denoted by (F) in fisheries models.</p> <p>Situation: Tony is about to go on two fishing trips in the up coming week. On Friday, he is going to Bear Lake, which is located near a factory that has been known to dump waste into the lake. On Saturday, he is going to Fox Lake, which is in a secluded valley.</p> <p>Q&A:Which day will Tony visit a lake that more likely has more fish? Saturday</p> <p>Q&A:Which lake probably has more fish in it? Fox Lake</p>	<p>Wrong Predicted Worlds</p> <p>Worlds: <i>[Bear Lake, Fox Lake]</i></p> <p>Effect: <i>fishing</i></p> <p>Cause^B: <i>Natural mortality</i></p> <p>Cause^S_{World1}: <i>dump waste</i></p> <p>Cause^S_{World2}: <i>secluded valley</i></p> <p>Cause Cmp: <i>World 1</i></p> <p>Relation: <i>negatively related</i></p> <p>Effect Cmp: <i>World 2</i></p> <p>Final Answer: <i>Fox Lake</i></p>
<p>ID:4215374242</p> <p>Background: Making these healthy lifestyle choices can also help prevent some types of cancer. In addition, you can lower the risk of cancer by avoiding carcinogens , which are substances that cause cancer. For example, you can reduce your risk of lung cancer by not smoking. You can reduce your risk of skin cancer by using sunscreen. How to choose a sunscreen that offers the most protection is explained below (Figure below). Some people think that tanning beds are a safe way to get a tan. This is a myth. Tanning beds expose the skin to UV radiation. Any exposure to UV radiation increases the risk of skin cancer. It doesn't matter whether the radiation comes from tanning lamps or the sun.</p> <p>Situation: Steve and Bill are really good friends with each other. The other day they were talking about some habits they have. Steve likes to work out and stay in shape, eats healthy and does not smoke. Bill said he wants to be more like Steve. Bill currently smokes, and loves to go out tanning in the sun without out sunscreen.</p> <p>Q&A:Who has a more likely chance to get lung cancer in the future? Bill</p>	<p>Wrong Predicted Comparison Results</p> <p>Worlds: <i>[Steve, Bill]</i></p> <p>Effect: <i>risk of lung cancer</i></p> <p>Cause^B: <i>not smoking.</i></p> <p>Cause^S_{World1}: <i>does not smoke</i></p> <p>Cause^S_{World2}: <i>smokes</i></p> <p>Cause Cmp: <i>World 2</i></p> <p>Relation: <i>negatively related</i></p> <p>Effect Cmp: <i>World 1</i></p> <p>Final Answer: <i>Steve</i></p>
<p>ID: 4035582237</p> <p>Background: Sometimes muscles and tendons get injured when a person starts doing an activity before they have warmed up properly. A warm up is a slow increase in the intensity of a physical activity that prepares muscles for an activity. Warming up increases the blood flow to the muscles and increases the heart rate. Warmed-up muscles and tendons are less likely to get injured. For example, before running or playing soccer, a person might jog slowly to warm muscles and increase their heart rate. Even elite athletes need to warm up (Figure below).</p> <p>Situation: Greg and Carl and about to do a marathon. Greg sees Carl doing some warm ups and laughs to himself and thinks it is silly. They both want to get a good time, and are both avid runners.</p> <p>Q&A:Who is more likely to get an injury during the race? Greg</p>	<p>Wrong Predicted Cause Property in Situation</p> <p>Worlds: <i>[Greg, Carl]</i></p> <p>Effect: <i>get injured.</i></p> <p>Cause^B: <i>Warmed-up</i></p> <p>Cause^S_{World1}: <i>warm ups</i></p> <p>Cause^S_{World2}: <i>warm ups</i></p> <p>Cause Cmp: <i>World 1</i></p> <p>Relation: <i>negatively related</i></p> <p>Effect Cmp: <i>World 2</i></p> <p>Final Answer: <i>Carl</i></p>

Table 9: Type 1 error cases made by our model

Examples	Prediction
<p>ID:710693196</p> <p>Background: Fish reproduce sexually. They lay eggs that can be fertilized either inside or outside of the body. In most fish, the eggs develop outside of the mother's body. In the majority of these species, fertilization also takes place outside the mother's body. The male and female fish release their gametes into the surrounding water, where fertilization occurs. Female fish release very high numbers of eggs to increase the chances of fertilization.</p> <p>Situation: All marine creatures are not fish. There are some mammals, for example, whales, also live in the water. Rob wants to know more about differences between fish and other non fish creatures in the water. He divided them into two groups, group A and group B. Group A consists of fish, and group B consists of non fish creatures in the water. He started to see the differences between these two groups</p> <p>Q&A:In group B, would fertilization most likely take place inside or outside of mother's body? inside</p>	<p>Missing Knowledge</p> <p>Worlds: <i>[Group A, group B]</i></p> <p>Effect: <i>fish</i></p> <p>Cause^B: <i>fertilization</i></p> <p>Cause^S_{World1}: <i>fish</i></p> <p>Cause^S_{World2}: <i>non fish creatures</i></p> <p>Cause Cmp: <i>World 1</i></p> <p>Relation: <i>positively related</i></p> <p>Effect Cmp: <i>World 1</i></p> <p>Final Answer: <i>Group A</i></p>
<p>ID:3339143431</p> <p>Background: In exercises such as weight lifting, skeletal muscle contracts against a resisting force (see Figure below). Using skeletal muscle in this way increases its size and strength. In exercises such as running, the cardiac muscle contracts faster and the heart pumps more blood. Using cardiac muscle in this way increases its strength and efficiency. Continued exercise is necessary to maintain bigger, stronger muscles. If you don't use a muscle, it will get smaller and weaker—so use it or lose it.</p> <p>Situation: A study was done in the town of Greenwich comparing muscle strength to the amount a person exercises. Mathew goes to the gym 5 times a week and does very intensive workouts. Damen on the other hand does not go to the gym at all and lives a mostly sedentary lifestyle.</p> <p>Q&A:Given Mathew suffers an injury while working out and cannot go to the gym for 3 months, will Mathews strength increase or decrease? decrease</p>	<p>Implicit Worlds</p> <p>Worlds: <i>[Mathew, Damen]</i></p> <p>Effect: <i>strength and efficiency</i></p> <p>Cause^B: <i>exercises such as running</i></p> <p>Cause^S_{World1}: <i>goes to the gym</i></p> <p>Cause^S_{World2}: <i>does not go to the gym</i></p> <p>Cause Cmp: <i>World 1</i></p> <p>Relation: <i>positively related</i></p> <p>Effect Cmp: <i>World 1</i></p> <p>Final Answer: <i>Mathew</i></p>
<p>ID: 2918297602</p> <p>Background: In common everyday speech, speed of sound refers to the speed of sound waves in air. However, the speed of sound varies from substance to substance: sound travels most slowly in gases; it travels faster in liquids; and faster still in solids. For example, (as noted above), sound travels at 343 m/s in air; it travels at 1,480 m/s in water (4.3 times as fast as in air); and at 5,120 m/s in iron (about 15 times as fast as in air). In an exceptionally stiff material such as diamond, sound travels at 12,000 metres per second (27,000 mph);[1] (about 35 times as fast as in air) which is around the maximum speed that sound will travel under normal conditions.</p> <p>Situation: John and Keith are neighbors. They have been pondering about how to communicate with each other in a doomsday scenario when all the electronic devices would be useless. They connected their houses with three ducts. One of the ducts is filled with air; they called it channel A. Another duct is filled with water; they called it channel B. And the last duct is filled with iron; they called it channel C. They can now transmit sound with these channels of communication; in case, disaster strikes</p> <p>Q&A:How much faster would be channel B than channel A in m/s, 1130 m/s or 1137 m/s?1137 m/s</p>	<p>Additional Math Computation</p> <p>Worlds: <i>[channel A,channel B]</i></p> <p>Effect: <i>5,120 m/s</i></p> <p>Cause^B: <i>water</i></p> <p>Cause^S_{World1}: <i>air</i></p> <p>Cause^S_{World2}: <i>water</i></p> <p>Cause Cmp: <i>World 1</i></p> <p>Relation: <i>positively related</i></p> <p>Effect Cmp: <i>World 1</i></p> <p>Final Answer: <i>channel A</i></p>

Table 10: Type 2 error cases could not be solved by our model

Examples	Prediction
<p>ID:1629973236</p> <p>Background:One result of air pollution is acid rain. Acid rain is precipitation with a low (acidic) pH. This rain can be very destructive to wildlife. When acid rain falls in forests, freshwater habitats, or soils, it can kill insects and aquatic life. It causes this damage because of its very low pH. Sulfur oxides and nitrogen oxides in the air both cause acid rain to form (Figure below). Sulfur oxides are chemicals that are released from coal-fired power plants. Nitrogen oxides are released from motor vehicle exhaust.</p> <p>Situation: Bill is planning on moving soon. He wants to move to a city that has fresher air and more wildlife to see. His two options that he must choose from are St. Louis and Seattle. Recently, Seattle has installed a new wind farm, and zero emission solar farm to generate power, while St. Louis recently installed a coal fired power plant. Both cities have similar commercial and industrial sectors, and only differ in how they generate power.</p> <p>Q&A: Which city will more likely have more vibrant wildlife? Seattle</p>	<p>Module Output:</p> <p>Worlds: <i>[Seattle,St. Louis]</i></p> <p>Effect: <i>wildlife</i></p> <p>Cause^B: <i>acid rain.</i></p> <p>Cause^S_{World1}: <i>wind farm,</i></p> <p>Cause^S_{World2}: <i>coal fired power plant.</i></p> <p>Cause Cmp: <i>World 2</i></p> <p>Relation: <i>negatively related</i></p> <p>Effect Cmp: <i>World 1</i></p> <p>Final Answer: <i>Seattle</i></p>
<p>ID:802123740</p> <p>Background:One result of air pollution is acid rain. Acid rain is precipitation with a low (acidic) pH. This rain can be very destructive to wildlife. When acid rain falls in forests, freshwater habitats, or soils, it can kill insects and aquatic life. It causes this damage because of its very low pH. Sulfur oxides and nitrogen oxides in the air both cause acid rain to form (Figure below). Sulfur oxides are chemicals that are released from coal-fired power plants. Nitrogen oxides are released from motor vehicle exhaust.</p> <p>Situation: Bill is planning on moving soon. He wants to move to a city that has fresher air and more wildlife to see. His two options that he must choose from are St. Louis and Seattle. Recently, Seattle has installed a new wind farm, and zero emission solar farm to generate power, while St. Louis recently installed a coal fired power plant. Both cities have similar commercial and industrial sectors, and only differ in how they generate power.</p> <p>Q&A: Will Seattle have more or less sulfur oxides in the air than St. Louis? less</p>	<p>Module Output:</p> <p>Worlds: <i>[Seattle,St. Louis]</i></p> <p>Effect: <i>Sulfur oxides</i></p> <p>Cause^B: <i>coal-fired power plants.</i></p> <p>Cause^S_{World1}: <i>wind farm,</i></p> <p>Cause^S_{World2}: <i>coal fired power plant.</i></p> <p>Cause Cmp: <i>World 2</i></p> <p>Relation: <i>positively related</i></p> <p>Effect Cmp: <i>World 2</i></p> <p>Final Answer: <i>St. Louis</i></p>
<p>ID:2133492859</p> <p>Background:Turner et al (2006) derived crash prediction models for this report’s predecessor and found a pronounced ‘2018safety in numbers’ effect in the models. Using the crash prediction model for mid-block locations, generic motorist and cyclist volumes can be used to demonstrate the impacts on the expected crash rate of varying motor vehicle and cycle volumes. As shown in figure 2.20, an increase in the proportion of cyclists to the overall traffic volume causes an increase in expected crashes at mid-block locations, but the crash rate increases at a decreasing rate. That is to say, the crash rate per cyclist goes down as the cycle volume increases.</p> <p>Situation: There were a lot of motorcycles on Interstate 17 last week. On Monday, there were 1355 motorcyclists. On Tuesday, there were 2355 motorcyclists. On Wednesday, there were 3351 motorcyclists. On Thursday, there were 4351 motorcyclists. On Friday, there were 5351 motorcyclists. On Saturday, there were 6351 motorcyclists. On Sunday, there were 7351 motorcyclists.</p> <p>Q&A: What day had a lower crash rate per cyclist: Thursday or Sunday? Sunday</p>	<p>Module Output:</p> <p>Worlds: <i>[Thursday,Sunday]</i></p> <p>Effect: <i>crash rate per cyclist</i></p> <p>Cause^B: <i>the cycle volume</i></p> <p>Cause^S_{World1}: <i>4351 motorcyclists</i></p> <p>Cause^S_{World2}: <i>7351 motorcyclists</i></p> <p>Cause Cmp: <i>World 2</i></p> <p>Relation: <i>negatively related</i></p> <p>Effect Cmp: <i>World 1</i></p> <p>Final Answer: <i>Thursday</i></p>
<p>ID:3099625752</p> <p>Background:The example of someone having a positive experience with a drug is easy to see how drug dependence and the law of effect works. The tolerance for a drug goes up as one continues to use it after having a positive experience with a certain amount the first time. It will take more and more to get that same feeling. This is when the controlled substance in an experiment would have to be modified and the experiment would really begin. The law of work for psychologist B. F. Skinner almost half a century later on the principles of operant conditioning, a learning process by which the effect, or consequence, of a response influences the future rate of production of that response.</p> <p>Situation: The Speed Squad met to discuss their experiences. They all said they always had a great experience using speed and used the same amount each time. They told how many times they used speed. Todd used it 36 times, Jesse used it 40 times, Craig used it 44 times, Alan used it 56 times, Shawn used it 69 times, Clarence used it 78 times, and Sean used it 86 times.</p> <p>Q&A: Who has a higher tolerance for speed: Alan or Clarence? Clarence</p>	<p>Module Output:</p> <p>Worlds: <i>[Alan,Clarence]</i></p> <p>Effect: <i>tolerance</i></p> <p>Cause^B: <i>positive experience with</i></p> <p>Cause^S_{World1}: <i>used it 56 times</i></p> <p>Cause^S_{World2}: <i>used it 78 times</i></p> <p>Cause Cmp: <i>World 2</i></p> <p>Relation: <i>positively related</i></p> <p>Effect Cmp: <i>World 2</i></p> <p>Final Answer: <i>Clarence</i></p>

Table 11: Examples correctly answered by our model in an interpretable manner.