# RECLOR: A READING COMPREHENSION DATASET REQUIRING LOGICAL REASONING

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Recent powerful pre-trained language models have achieved remarkable performance on most of the popular datasets for reading comprehension. It is time to introduce more challenging datasets to push the development of this field towards more comprehensive reasoning of text. In this paper, we introduce a new **Re**ading **C**omprehension dataset requiring **l**ogical **r**easoning (ReClor) extracted from standardized graduate admission examinations. As earlier studies suggest, human-annotated datasets usually contain biases, which are often exploited by models to achieve high accuracy without truly understanding the text. In order to comprehensively evaluate the logical reasoning ability of models on ReClor, we propose to identify biased data points and separate them into EASY set while the rest as HARD set. Empirical results show that the state-of-the-art models have an outstanding ability to capture biases contained in the dataset with high accuracy on EASY set. However, they struggle on HARD set with poor performance near that of random guess, indicating more research is needed to essentially enhance the logical reasoning ability of current models.

## 1 INTRODUCTION

Machine reading comprehension (MRC) is a fundamental task in Natural Language Processing (NLP), which requires models to understand a body of text and answer a particular question related to the context. With success of unsupervised representation learning in NLP, language pre-training based models such as GPT-2 (Radford et al., 2019), BERT (Devlin et al., 2019), XLNet (Yang et al., 2019) and RoBERTa (Liu et al., 2019) have achieved nearly saturated performance on most of the popular MRC datasets (Rajpurkar et al., 2016; Lai et al., 2017; Rajpurkar et al., 2018; Wang et al., 2018). It is time to challenge state-of-the-art models with more difficult reading comprehension tasks and move a step forward to more comprehensive analysis and reasoning over text (Dua et al., 2019).

In natural language understanding, logical reasoning is an important ability to examine, analyze and critically evaluate arguments as they occur in ordinary language (Council, 2019a). It is a significant component of human intelligence and is essential in negotiation, debate and writing *etc.* However, existing reading comprehension datasets have none or merely a small amount of data requiring logical reasoning, *e.g.*, 0% in MCTest dataset (Richardson et al., 2013) and 1.2% in SQuAD (Rajpurkar et al., 2016) according to Sugawara & Aizawa (2016). One related task is natural language inference, which requires models to label the logical relationships of sentence pairs. However, this task only considers three types of simple logical relationships and needs reasoning at sentence-level. To push the development of models in logical reasoning from simple logical relationship classification to multiple complicated logical reasoning and from sentence-level to passage-level, it is necessary to introduce a reading comprehension dataset targeting logical reasoning.

A typical example of logical reasoning questions is shown in Table 1. Similar to the format of multiple-choice reading comprehension datasets (Richardson et al., 2013; Lai et al., 2017), it contains a context, a question and four options with only one right answer. To answer the question in this example, readers need to identify the logical connections between the lines to pinpoint the conflict, then understand each of the options and select an option that solves the conflict. Human minds need extensive training and practice to get used to complex reasoning, and it will take immense efforts for crowdsourcing workers to design such logical reasoning questions. Inspired by the datasets extracted from standardized examinations (Lai et al., 2017; Clark et al., 2018), we build a dataset by selecting such logical reasoning questions from standardized exams such as GMAT [1] and LSAT [2]. We finally collect 6,139 pieces of logical reasoning questions, which constitute a **Re**ading **C**omprehension dataset requiring **lo**gical **r**easoning (**ReClor**).

Human-annotated datasets usually contain biases (Schwartz et al., 2017; Cai et al., 2017; Bugert et al., 2017; Poliak et al., 2018; Gururangan et al., 2018; Zellers et al., 2019), which are often exploited by neural network models as shortcut solutions to achieve high testing accuracy. For data points whose options can be selected correctly without knowing the contexts and questions, we classify them as biased ones. In order to fully assess the logical reasoning ability of the models, we propose to identify the biased data points and group them as EASY set, and put the rest into HARD set. Based on our experiments on these separate sets, we find that even the state-of-the-art models can only perform well on EASY set and struggle on HARD set as shown in Figure 1. This phenomenon shows that current models can well capture the biases in the dataset but lack the ability to understand the text and reason based on connections between the lines. On the other hand, human beings perform similarly on both the EASY and HARD set. It is thus observed that there is still a long way to go to equip models with true logical reasoning ability.

The contributions of our paper are two-fold. First, we introduce ReClor, a new reading comprehension dataset ReClor requiring logical reasoning. We use option-only-input baselines trained with different random seeds to identify the data points with biases in testing set, and group them as EASY set, with the rest as HARD set to facilitate comprehensive evaluation. Second, we evaluate several state-of-the-art models on ReClor and find these pre-trained language models can perform well on EASY set but struggle on the HARD set. This indicates although current models are good at exploiting biases in the dataset, they are far from capable of performing real logical reasoning yet.

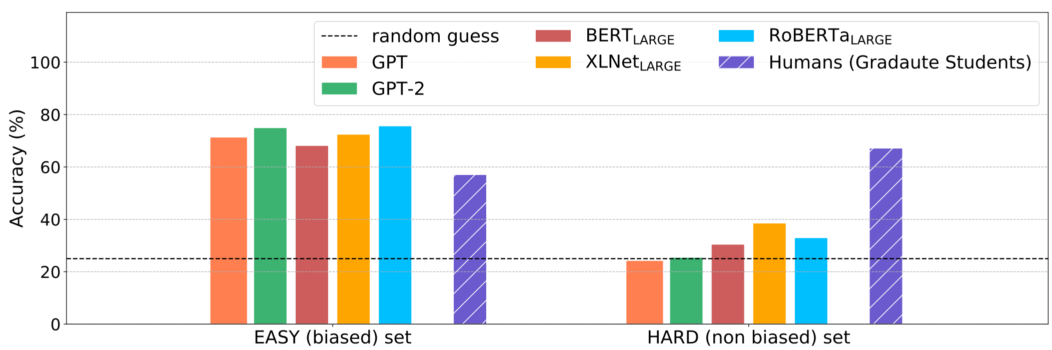


Figure 1: Performance comparison of state-of-the-art models and humans (graduate students) on EASY and HARD set of ReClor testing set.

## 2 Related Work

**Reading Comprehension Datasets** A variety of reading comprehension datasets have been introduced to promote the development of this field. MCTest (Richardson et al., 2013) is a dataset with 2,000 multiple-choice reading comprehension questions about fictional stories in the format similar to ReClor. Rajpurkar et al. (2016) proposed SQuAD dataset, which contains 107,785 question-answer pairs on 536 Wikipedia articles. The authors manually labeled 192 examples of the dataset and found that the examples mainly require reasoning of lexical or syntactic variation. In an analysis of the above-mentioned datasets, Sugawara & Aizawa (2016) found that none of questions requiring logical reasoning in MCTest dataset (Richardson et al., 2013) and only 1.2% in SQuAD dataset (Rajpurkar et al., 2016). Lai et al. (2017) introduced RACE dataset by collecting the English exams for middle and high school Chinese students in the age range between 12 to 18. They hired crowd worker on Amazon Mechanical Turk to label the reasoning type of 500 samples in the dataset and

[1] https://en.wikipedia.org/wiki/Graduate_Management_Admission_Test
[2] https://en.wikipedia.org/wiki/Law_School_Admission_Test

| |
|---|
| **Context:**<br>In jurisdictions where use of headlights is optional when visibility is good, drivers who use headlights at all times are less likely to be involved in a collision than are drivers who use headlights only when visibility is poor. Yet Highway Safety Department records show that making use of headlights mandatory at all times does nothing to reduce the overall number of collisions.<br>**Question:**<br>Which one of the following, if true, most helps to resolve the apparent discrepancy in the information above?<br>**Options:**<br>A. In jurisdictions where use of headlights is optional when visibility is good, one driver in four uses headlights for daytime driving in good weather.<br>B. Only very careful drivers use headlights when their use is not legally required.<br>C. The jurisdictions where use of headlights is mandatory at all times are those where daytime visibility is frequently poor.<br>D. A law making use of headlights mandatory at all times is not especially difficult to enforce.<br>**Answer:** B |

Table 1: An example from the ReClor dataset which is modified from the Law School Admission Council (2019b).

show that around 70 % of the samples are in the category of word matching, paraphrasing or single-sentence reasoning. To encourage progress on deeper comprehension of language, more reading comprehension datasets requiring more complicated reasoning types are introduced, such as multi-hop reasoning across multiple sentences (Khashabi et al., 2018) and multiple documents (Welbl et al., 2018), iterative reasoning about the narrative of a story (Kočiskỳ et al., 2018), commonsense knowledge reasoning (Mihaylov et al., 2018; Zhang et al., 2018) and numerical discrete reasoning over paragraphs (Dua et al., 2019). However, to the best of our knowledge, although there are some datasets targeting logical reasoning in other NLP tasks mentioned in the next section, there is no dataset targeting evaluating logical reasoning in reading comprehension task. This work introduces a new dataset to fill this gap.

**Logical Reasoning in NLP** There are several tasks and datasets introduced to investigate logical reasoning in NLP. The task of *natural language inference*, also known as *recognizing textual entailment* (Fyodorov et al., 2000; Condoravdi et al., 2003; Bos & Markert, 2005; Dagan et al., 2005; MacCartney & Manning, 2009) requires models to take a pair of sentence as input and classify their relationship types, *i.e.*, ENTAILMENT, NEUTRAL, or CONTRADICTION. SNLI (Bowman et al., 2015) and MultiNLI (Williams et al., 2018) datasets are proposed for this task. However, this task only focuses on sentence-level logical relationship reasoning and the relationships are limited to only a few types. Another task related to logical reasoning in NLP is *argument reasoning comprehension task* introduced by Habernal et al. (2018) with a dataset of this task. Given an argument with a claim and a premise, this task aims to select the correct implicit warrant from two options. Although the task is on passage-level logical reasoning, it is limited to only one logical reasoning type, *i.e.*, identifying warrants. ReClor and the proposed task integrate various logical reasoning types into reading comprehension, with the aim to promote the development of models in logical reasoning not only from sentence-level to passage-level, but also from simple logical reasoning types to the complicated diverse ones.

**Datasets from Examinations** There have been several datasets extracted from human standardized examinations in NLP, such as RACE dataset (Lai et al., 2017) mentioned above. Besides RACE, NTCIR QA Lab (Shibuki et al., 2014) offers comparative evaluation for solving real-world university entrance exam questions; The dataset of CLEF QA Entrance Exams Task (Rodrigo et al., 2015) is extracted from standardized English examinations for university admission in Japan; ARC dataset (Clark et al., 2018) consists of 7,787 science questions targeting student grade level, ranging from 3rd grade to 9th. Compared with these datasets, ReClor distinguishes itself by targeting logical reasoning.

| TYPE | DESCRIPTION |
|---|---|
| Necessary Assumptions (11.5%) | identify the claim that must be true or is required in order for the argument to work. |
| Sufficient Assumptions (5.9%) | identify a sufficient assumption, that is, an assumption that, if added to the argument, would make it logically valid. |
| Strengthen (9.3%) | identify information that would strengthen an argument |
| Weaken (11.4%) | identify information that would weaken an argument |
| Evaluation (1.3%) | identify information that would be useful to know to evaluate an argument |
| Conclusion (3.1%) | identify the conclusion of a line of reasoning |
| Summary/Main Point (1.3%) | identify the main point of an argument |
| Implication (4.4%) | identify something that follows logically from a set of premises |
| Most Strongly Supported (5.8%) | find the choice that is most strongly supported by a stimulus |
| Explain or Resolve (8.6%) | identify information that would explain or resolve a situation |
| Match Principles/Points (2.6%) | match principles from the stimulus to principles in a choice |
| Match Flaws (2.9%) | find a choice containing an argument that exhibits the same flaws as the passages argument |
| Match the Structure (2.8%) | match the structure of an argument in a choice to the structure of the argument in the passage |
| Issue (3.1%) | identify or infer an issue in dispute |
| Technique (3.1%) | identify the technique used in the reasoning of an argument |
| Role (3.2%) | describe the individual role that a statement is playing in a larger argument |
| Identify a Flaw (11.5%) | identify a flaw in an arguments reasoning |
| Others (8.2%) | other types of questions which are not included by the above |

Table 2: The percentage and description of each logical reasoning type. The descriptions are adapted from those specified by Khan Academy (2019).

## 3 ReClor Data Collection and Analysis

### 3.1 Data collection

The format of data in ReClor is similar to other multiple-choice reading comprehension datasets (Richardson et al., 2013; Lai et al., 2017), where a data point contains a context, a question and four answer options, among which only one option is right/most suitable. We collect reading comprehension problems that require complicated logical reasoning. However, producing such data requires the ability to perform complex logical reasoning, which makes it hard for crowdsourcing workers to generate such logical questions. Fortunately, we find the reading comprehension problems in some standardized tests, such as GMAT and LSAT, are highly in line with our expectation.

We construct a dataset containing 6,139 logical reasoning questions sourced from open websites and books. In the original problems, there are five answer options in which only one is right. To comply with fair use of law [3], we shuffle the order of answer options and randomly delete one of the wrong options for each data point, which results in four options with one right option and three wrong options. Furthermore, similar to ImageNet dataset [4], we plan to offer the dataset to researchers/educators who agree to have it for non-commercial research and/or educational use only.

### 3.2 Types of Logical Reasoning Questions

As mentioned above, we collect 6,139 data points, which are divided into training set, validation set and testing set with 4,651, 500 and 1,000 data points respectively. We analyze and manually annotate the types of questions on testing set and group them into 18 categories, whose percentages and descriptions are shown in Table 2. Examples of different types of logical reasoning are listed in the appendix.

[3] `https://www.copyright.gov/fair-use/more-info.html`
[4] `http://image-net.org/download-faq`

### 3.3 DATA BIASES IN THE DATASET

The dataset is collected from exams devised by experts in logical reasoning, which means it is annotated by humans and may introduce biases in the dataset. Recent studies have shown that models can utilize the biases in a dataset of natural language understanding to perform well on the task without truly understanding the text (Schwartz et al., 2017; Cai et al., 2017; Bugert et al., 2017; Poliak et al., 2018; Gururangan et al., 2018; Zellers et al., 2019). It is necessary to analyze such data biases to help evaluate models. In the ReClor dataset, the common context and question are shared across the four options for each data point, so we focus on the analysis of the difference in lexical choice and sentence length of the right and wrong options without contexts and questions. We first investigate the biases of lexical choice. We lowercase the options and then use WordPiece tokenization (Wu et al., 2016) of $\text{BERT}_{\text{BASE}}$ (Devlin et al., 2019) to get the tokens. Similar to Poliak et al. (2018), for the tokens in options, we analyze their conditional probability of label $l \in \{\text{right}, \text{wrong}\}$ given by the token $t$ by

$$p(l|t) = \frac{count(t,l)}{count(t)}. \quad (1)$$

The larger the above correlation score is for a particular token, the more likely it contributes to the prediction of related option. Table 3 reports tokens which occur at least twenty times with the highest scores since many of the tokens with the highest scores are of low frequency. We further analyze the lengths of right and wrong options (Gururangan et al., 2018). We notice a slight difference in the distribution of sentence length for right and wrong options. The average length for wrong options is around 21.78 whereas that for right options is generally longer with an average length of 23.02.

| TOKEN | SCORE | FREQ |
|---|---|---|
| motive | 65.2 | 23 |
| ##ce | 60.7 | 28 |
| thereby | 56.0 | 25 |
| consequence | 52.4 | 21 |
| warm | 52.4 | 21 |
| interfere | 52.2 | 23 |
| contributes | 52.2 | 23 |
| manufacture | 52.0 | 25 |
| included | 52.0 | 25 |
| preferences | 52.0 | 25 |

Table 3: Top 10 tokens that correlate to right options with more than 20 occurrences.

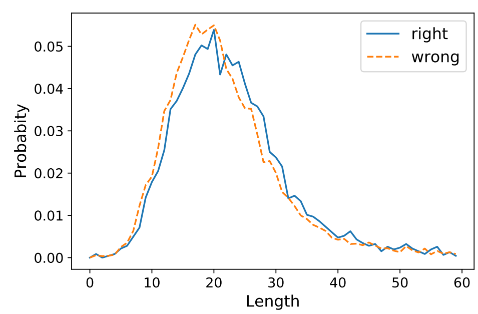


Figure 2: The distribution of the option length in ReClor with respect to right and wrong labels.

## 4 BASELINE MODELS

Many neural network based models have achieved impressive results in various NLP tasks. We challenge these neural models with ReClor to investigate how well they can perform. The input format of different models are shown in Table 4.

**fastText** FastText (Joulin et al., 2017) models sentences as a bag of n-grams, and tries to predict the probability of each answer being correct independently. We choose the answer with the highest score as the prediction for the multiple-choice setting.

**LSTM sentence encoder** A two-layer bi-LSTM is randomly initialized as a sentence encoder with GloVe word embedding (Pennington et al., 2014). With a span of text as input, the last hidden state of the second layer is max-pooled and then fed into a fully-connected layer to compute the output score.

**GPT and GPT-2** GPT (Radford et al., 2018) and GPT-2 (Radford et al., 2019) are both transformer (Vaswani et al., 2017) based models which are pre-trained using unsupervised method with a standard language modeling objective. GPT is pre-trained on BooksCorpus; GPT-2 is pre-trained using a larger dataset called WebText. Here we use the smallest model proposed in (Radford et al., 2019) as our GPT-2 baseline. To fine-tune on ReClor, the final hidden vector corresponding to the last input token (`[_classify_]`) is used as the aggregate representation followed by an extra fully connected layer to compute the score.

<table>
<tr><th>MODEL</th><th>INPUT FORMAT</th></tr>
<tr><td>GPT Radford et al. (2018)</td><td>_start_ Context _delimiter_ Question || Option _classify_</td></tr>
<tr><td>GPT-2 Radford et al. (2019)</td><td>_start_ Context _delimiter_ Question || Option _classify_</td></tr>
<tr><td>BERT (Devlin et al., 2019)</td><td>[CLS] Context [SEP] Question || Option [SEP] [PAD]...</td></tr>
<tr><td>XLNet (Yang et al., 2019)</td><td><pad>... Context <sep> Question || Option <sep> <cls></td></tr>
<tr><td>RoBERTa (Liu et al., 2019)</td><td><s> Context </s> </s> Question || Option </s> <pad>...</td></tr>
</table>

Table 4: Input formats of different models. Context, Question and Option represent the token sequences of the context, question and option respectively, and || denotes concatenation.

**BERT** BERT (Devlin et al., 2019) is also a transformer (Vaswani et al., 2017) based model which is trained by using BooksCorpus (Zhu et al., 2015) and English Wikipedia in two unsupervised tasks, i.e., Masked LM (MLM) and Next Sentence Prediction (NSP). During fine-tuning, the final hidden vector corresponding to the first input token ([CLS]) is used as the aggregate representation followed by an extra fully connected layer to compute the score.

**XLNet** XLNet (Yang et al., 2019) is trained with Permutation Language Modeling and without NSP. In addition, beside BooksCorpus and English Wikipedia used in BERT, it uses Giga5 (Parker et al., 2011), ClueWeb 2012-B (extended from (Callan et al., 2009)), and Common Crawl (com, 2019) for pre-training. We use the final hidden vector corresponding to the last input token <cls> as the aggregate representation and introduce two fully connected layers to predict the score.

**RoBERTa** RoBERTa (Liu et al., 2019) is an improved pre-training procedure of BERT with training the model longer, with bigger batches over more data and removing NSP objective *etc.*. Extra two fully connected layers are added to transform the final hidden vector of the first input token (<s> to the score.

# 5 EXPERIMENTS

## 5.1 EXPERIMENT DETAILS

Adam is used by all models with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 1e-8$. For fastText, we use its python library[5] by converting ReClor to the required form, and keep the default setting of the hyper parameters. For Bi-LSTM, we use a two-layer Bidirectional LSTM with the GloVe 300d word embedding (Pennington et al., 2014) followed by max-pooling and a fully-connected layer. We train the model for 100 epochs using a batch size of 64 and learning rate of 0.1. A learning rate decay of 0.5 is also applied every 10 epochs. For pre-training models, we modify code of PyTorch-Transformers of Hugging Face[6] to implement them on ReClor. We use a batch size of 24 and fine-tune for 50 epochs. For different models, we select the best fine-tuning learning rate among 6.25e-05, 2e-05, 1e-05, and 5e-06 on the validation set. The maximum input sequence length for all models is 256. More hyperparameters are shown in Table 7 in the appendix.

## 5.2 EXPERIMENTS TO FIND BIASED DATA

As mentioned earlier, biases prevalently exist in human-annotated datasets (Poliak et al., 2018; Gururangan et al., 2018; Zellers et al., 2019; Niven & Kao, 2019), which are often exploited by models to perform well without truly understanding the text. Therefore, it is necessary to find out the biased data points in ReClor in order to evaluate models in a more comprehensive manner. To this end, we feed the five strong baseline models (GPT, GPT-2, $\text{BERT}_{\text{BASE}}$, $\text{XLNet}_{\text{BASE}}$ and $\text{RoBERTa}_{\text{BASE}}$) with **ONLY THE ANSWER OPTIONS** for each problem. In other words, we purposely remove the context and question in the inputs. In this way, we are able to identify those problems that can be answered correctly by merely exploiting the biases in answer options without knowing the relevant context and question. However, the setting of this task is an MCQ question with 4 probable options, and even a chance baseline could have 25% probability to get it right. To eliminate the effect of random guess, we set four different random seeds for each model and pick the data points that are predicted correctly in all four cases to form the EASY set. Then, the data points which are predicted correctly by the models at random could be nearly eliminated, since any data point only

[5]https://github.com/facebookresearch/fastText

[6]https://github.com/huggingface/pytorch-transformers

| MODEL | VAL | TEST | NUMBER |
|---|---|---|---|
| Chance | 25.0 | 25.0 | 0.39 |
| GPT | 45.8 | 42.2 | 238 |
| GPT-2 | 46.8 | 42.6 | 245 |
| $\text{BERT}_{\text{BASE}}$ | 47.2 | 43.2 | 234 |
| $\text{XLNet}_{\text{BASE}}$ | 47.5 | 43.2 | 225 |
| $\text{RoBERTa}_{\text{BASE}}$ | 48.8 | 41.7 | 200 |
| Union | – | – | 440 |

Table 5: Average accuracy of each model using four different random seeds with only answer options as input, and the number of their common correct predictions.

has a probability of $(25\%)^4 = 0.39\%$ to be guessed right consecutively for four times. Then we unite the sets of data points that are consistently predicted right by each model, because intuitively different models may learn different biases of the dataset. The above process is formulated as the following expression,

$$
\begin{aligned}
\mathbb{C}_{\mathrm{EASY}} = &(\mathbb{C}^{\mathrm{seed}_1}_{\mathrm{GPT}} \cap \mathbb{C}^{\mathrm{seed}_2}_{\mathrm{GPT}} \cap \mathbb{C}^{\mathrm{seed}_3}_{\mathrm{GPT}} \cap \mathbb{C}^{\mathrm{seed}_4}_{\mathrm{GPT}}) \\
&\cup (\mathbb{C}^{\mathrm{seed}_1}_{\mathrm{GPT-2}} \cap \mathbb{C}^{\mathrm{seed}_2}_{\mathrm{GPT-2}} \cap \mathbb{C}^{\mathrm{seed}_3}_{\mathrm{GPT-2}} \cap \mathbb{C}^{\mathrm{seed}_4}_{\mathrm{GPT-2}}) \\
&\cup (\mathbb{C}^{\mathrm{seed}_1}_{\mathrm{BERT}} \cap \mathbb{C}^{\mathrm{seed}_2}_{\mathrm{BERT}} \cap \mathbb{C}^{\mathrm{seed}_3}_{\mathrm{BERT}} \cap \mathbb{C}^{\mathrm{seed}_4}_{\mathrm{BERT}}) \\
&\cup (\mathbb{C}^{\mathrm{seed}_1}_{\mathrm{XLNet}} \cap \mathbb{C}^{\mathrm{seed}_2}_{\mathrm{XLNet}} \cap \mathbb{C}^{\mathrm{seed}_3}_{\mathrm{XLNet}} \cap \mathbb{C}^{\mathrm{seed}_4}_{\mathrm{XLNet}}) \\
&\cup (\mathbb{C}^{\mathrm{seed}_1}_{\mathrm{RoBERTa}} \cap \mathbb{C}^{\mathrm{seed}_2}_{\mathrm{RoBERTa}} \cap \mathbb{C}^{\mathrm{seed}_3}_{\mathrm{RoBERTa}} \cap \mathbb{C}^{\mathrm{seed}_4}_{\mathrm{RoBERTa}}),
\end{aligned} \tag{2}
$$

$$\mathbb{C}_{\mathrm{HARD}} = \mathbb{C}_{\mathrm{TEST}} - \mathbb{C}_{\mathrm{EASY}},$$

where $\mathbb{C}^{\mathrm{seed}_1}_{\mathrm{BERT}}$ denotes the set of data points which are predicted correctly by $\text{BERT}_{\text{BASE}}$ with seed 1, and similarly for the rest. Table 5 shows the average performance for each model trained with four different random seeds and the number of data points predicted correctly by all of them. Finally, we get 440 data points from testing set $\mathbb{C}_{\mathrm{TEST}}$ and we denote this subset as EASY set $\mathbb{C}_{\mathrm{EASY}}$ and the other as HARD set $\mathbb{C}_{\mathrm{HARD}}$.

### 5.3 Results and Analysis

The performance of all tested models on the ReClor is presented in Table 6. This dataset is built on questions designed for students who apply for admission to graduate schools, thus we randomly choose 100 samples from testing set and divide them into ten tests, which are distributed to ten different graduate students in a university. We take the average of their scores and present it as the baseline of graduate students. The data of ReClor are carefully chosen and modified from only high quality questions from standardized graduate entrance exams. We set the ceiling performance to 100% since ambiguous questions are not included in the dataset.

The performance of fastText is better than random guess, showing that word correlation could be used to help improve performance to some extent. It is difficult for Bi-LSTM to converge on this dataset. Transformer-based pre-training models have relatively good performance, close to the performance of graduate students. However, we find that these models only perform well on EASY set with around 70% accuracy, showing these models have an outstanding ability to capture the biases of the dataset, but they perform poorly on HARD set with only around 30% accuracy. In contrast, humans can still keep good performance on HARD set. We notice the difference in testing accuracy performed by graduate students on EASY and HARD set, but this could be due to the small number of students participated in the experiments. Therefore, we say humans perform relatively consistent on both biased and non-biased dataset.

We further analyze the model performance with respect to different question types of logical reasoning. Some results are shown in Figure 3 and the full results are shown in Figure 4, 5 and 6 in the appendix. Three models of $\text{BERT}_{\text{LARGE}}$, $\text{XLNetT}_{\text{LARGE}}$ and $\text{RoBERTa}_{\text{LARGE}}$ all perform poorly on the type of Match Principles/Points on EASY set. On HARD set, the three models perform

| MODEL | VAL | TEST | TEST-EASY | TEST-HARD |
|---|---|---|---|---|
| Chance | 25.0 | 25.0 | 25.0 | 25.0 |
| fastText | 32.0 | 30.8 | 40.2 | 23.4 |
| Bi-LSTM | 29.0 | 25.5 | 25.2 | 25.7 |
| GPT | 51.0 | 45.0 | 71.4 | 24.3 |
| GPT-2 | 52.2 | 47.3 | 75.0 | 25.5 |
| $\text{BERT}_{\text{BASE}}$ | 50.2 | 44.5 | 69.5 | 24.8 |
| $\text{BERT}_{\text{LARGE}}$ | 51.6 | 47.1 | 68.2 | 30.5 |
| $\text{XLNet}_{\text{BASE}}$ | 54.8 | 48.9 | 75.5 | 28.0 |
| $\text{XLNetT}_{\text{LARGE}}$ | 59.4 | 53.5 | 72.5 | 38.6 |
| $\text{RoBERTa}_{\text{BASE}}$ | 56.0 | 48.1 | 70.7 | 30.4 |
| $\text{RoBERTa}_{\text{LARGE}}$ | 57.8 | 51.8 | 75.7 | 33.0 |
| Graduate Students | – | 63.0 | 57.1 | 67.2 |
| Ceiling Performance | – | 100 | 100 | 100 |

Table 6: Accuracy (%) of models and human performance.

poorly on certain types such as STRENGTHEN, WEAKEN and ROLE which require extensive logical reasoning. However, they perform relatively better on other types, such as CONCLUSION and SUMMARY/MAIN POINT that are more straight-forward.

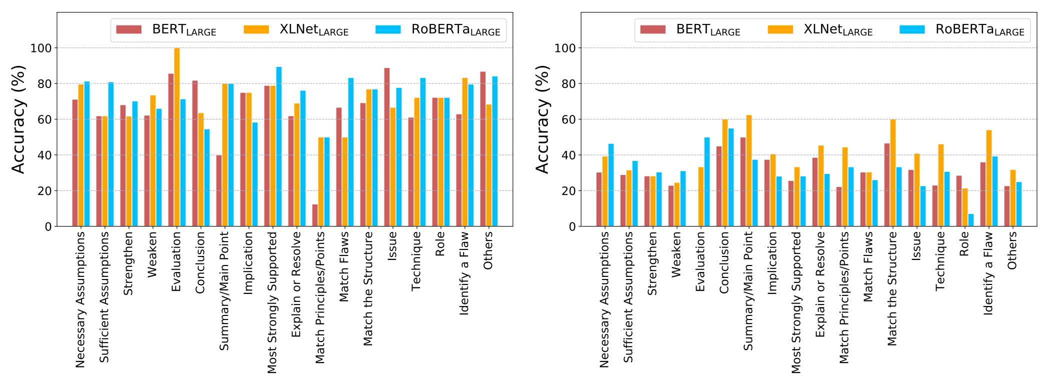


Figure 3: Performance of models on EASY (left) and HARD (right) testing sets.

## 6 CONCLUSION

In this paper, we introduce ReClor, a reading comprehension dataset requiring logical reasoning, with the aim to push research progress on logical reasoning in NLP forward from sentence-level to passage-level and from simple logical reasoning to multiple complicated one. We propose to identify biased data points and split the testing dataset into EASY and HARD group for biased and non-biased data separately. We further empirically study the different behaviors of state-of-the-art models on these two testing sets, and find recent powerful transformer-based pre-trained language models have an excellent ability to exploit the biases in the dataset but have difficulty in understanding and reasoning given the non-biased data with low performance close to or slightly better than random guess. These results show there is a long way to equip deep learning models with real logical reasoning abilities. We hope this work would inspire more research in future to adopt similar split technique and evaluation scheme when reporting their model performance.

## REFERENCES

Common crawl. http://http://commoncrawl.org, 2019.

Khan Academy. https://www.khanacademy.org/test-prep/lsat/lsat-lessons/logical-reasoning/a/logical-reasoning--article--question-type-catalog, 2019. Accessed Sept. 16, 2019.

Johan Bos and Katja Markert. Recognising textual entailment with logical inference. In *Proceedings of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*, pp. 628–635. Association for Computational Linguistics, 2005.

Samuel R Bowman, Gabor Angeli, Christopher Potts, and Christopher D Manning. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 632–642, 2015.

Michael Bugert, Yevgeniy Puzikov, Andreas Rücklé, Judith Eckle-Kohler, Teresa Martin, Eugenio Martínez-Cámara, Daniil Sorokin, Maxime Peyrard, and Iryna Gurevych. Lsdsem 2017: Exploring data generation methods for the story cloze test. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pp. 56–61, 2017.

Zheng Cai, Lifu Tu, and Kevin Gimpel. Pay attention to the ending: Strong neural baselines for the roc story cloze task. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 616–622, 2017.

Jamie Callan, Mark Hoy, Changkuk Yoo, and Le Zhao. Clueweb09 data set, 2009.

Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.

Cleo Condoravdi, Dick Crouch, Valeria De Paiva, Reinhard Stolle, and Daniel G Bobrow. Entailment, intensionality and text understanding. In *Proceedings of the HLT-NAACL 2003 workshop on Text meaning*, pp. 38–45, 2003.

Law School Admission Council. https://www.lsac.org/lsat/taking-lsat/test-format/logical-reasoning, 2019a. Accessed Sept. 16, 2019.

Law School Admission Council. https://www.lsac.org/lsat/taking-lsat/test-format/logical-reasoning/logical-reasoning-sample-questions, 2019b. Accessed Sept. 16, 2019.

Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pp. 177–190. Springer, 2005.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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)*, pp. 4171–4186, 2019.

Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. In *Proceedings of NAACL-HLT*, pp. 2368–2378, 2019.

Yaroslav Fyodorov, Yoad Winter, and Nissim Francez. A natural logic inference system. In *Proceedings of the 2nd Workshop on Inference in Computational Semantics (ICoS-2)*. Citeseer, 2000.

Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. Annotation artifacts in natural language inference data. *arXiv preprint arXiv:1803.02324*, 2018.

Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. The argument reasoning comprehension task: Identification and reconstruction of implicit warrants. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1930–1940, 2018.

Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. Bag of tricks for efficient text classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 427–431. Association for Computational Linguistics, April 2017.

Daniel Khashabi, Snigdha Chaturvedi, Michael Roth, Shyam Upadhyay, and Dan Roth. Looking beyond the surface: A challenge set for reading comprehension over multiple sentences. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 252–262, 2018.

Tomáš Kočiský, Jonathan Schwarz, Phil Blunsom, Chris Dyer, Karl Moritz Hermann, Gábor Melis, and Edward Grefenstette. The narrativeqa reading comprehension challenge. *Transactions of the Association for Computational Linguistics*, 6:317–328, 2018.

Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. Race: Large-scale reading comprehension dataset from examinations. *arXiv preprint arXiv:1704.04683*, 2017.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Bill MacCartney and Christopher D Manning. An extended model of natural logic. In *Proceedings of the eighth international conference on computational semantics*, pp. 140–156. Association for Computational Linguistics, 2009.

Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.

Timothy Niven and Hung-Yu Kao. Probing neural network comprehension of natural language arguments. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4658–4664, 2019.

Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. English gigaword fifth edition, 2011.

Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.

Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. Hypothesis only baselines in natural language inference. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pp. 180–191, 2018.

Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *URL https://s3-us-west-2. amazonaws. com/openai-assets/researchcovers/languageunsupervised/language understanding paper. pdf*, 2018.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8), 2019.

Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.

Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions for squad. *arXiv preprint arXiv:1806.03822*, 2018.

Matthew Richardson, Christopher JC Burges, and Erin Renshaw. Mctest: A challenge dataset for the open-domain machine comprehension of text. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 193–203, 2013.

Alvaro Rodrigo, Anselmo Penas, Yusuke Miyao, Eduard H Hovy, and Noriko Kando. Overview of clef qa entrance exams task 2015. In *CLEF (Working Notes)*, 2015.

Roy Schwartz, Maarten Sap, Ioannis Konstas, Leila Zilles, Yejin Choi, and Noah A Smith. Story cloze task: Uw nlp system. In *Proceedings of the 2nd Workshop on Linking Models of Lexical, Sentential and Discourse-level Semantics*, pp. 52–55, 2017.

Hideyuki Shibuki, Kotaro Sakamoto, Yoshinobu Kano, Teruko Mitamura, Madoka Ishioroshi, Kelly Y Itakura, Di Wang, Tatsunori Mori, and Noriko Kando. Overview of the ntcir-11 qa-lab task. In *Ntcir*, 2014.

Saku Sugawara and Akiko Aizawa. An analysis of prerequisite skills for reading comprehension. In *Proceedings of the Workshop on Uphill Battles in Language Processing: Scaling Early Achievements to Robust Methods*, pp. 1–5, 2016.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355, 2018.

Johannes Welbl, Pontus Stenetorp, and Sebastian Riedel. Constructing datasets for multi-hop reading comprehension across documents. *Transactions of the Association for Computational Linguistics*, 6:287–302, 2018.

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

Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.

Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *arXiv preprint arXiv:1906.08237*, 2019.

Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.

Sheng Zhang, Xiaodong Liu, Jingjing Liu, Jianfeng Gao, Kevin Duh, and Benjamin Van Durme. Record: Bridging the gap between human and machine commonsense reading comprehension. *arXiv preprint arXiv:1810.12885*, 2018.

Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pp. 19–27, 2015.

# 7 Appendix

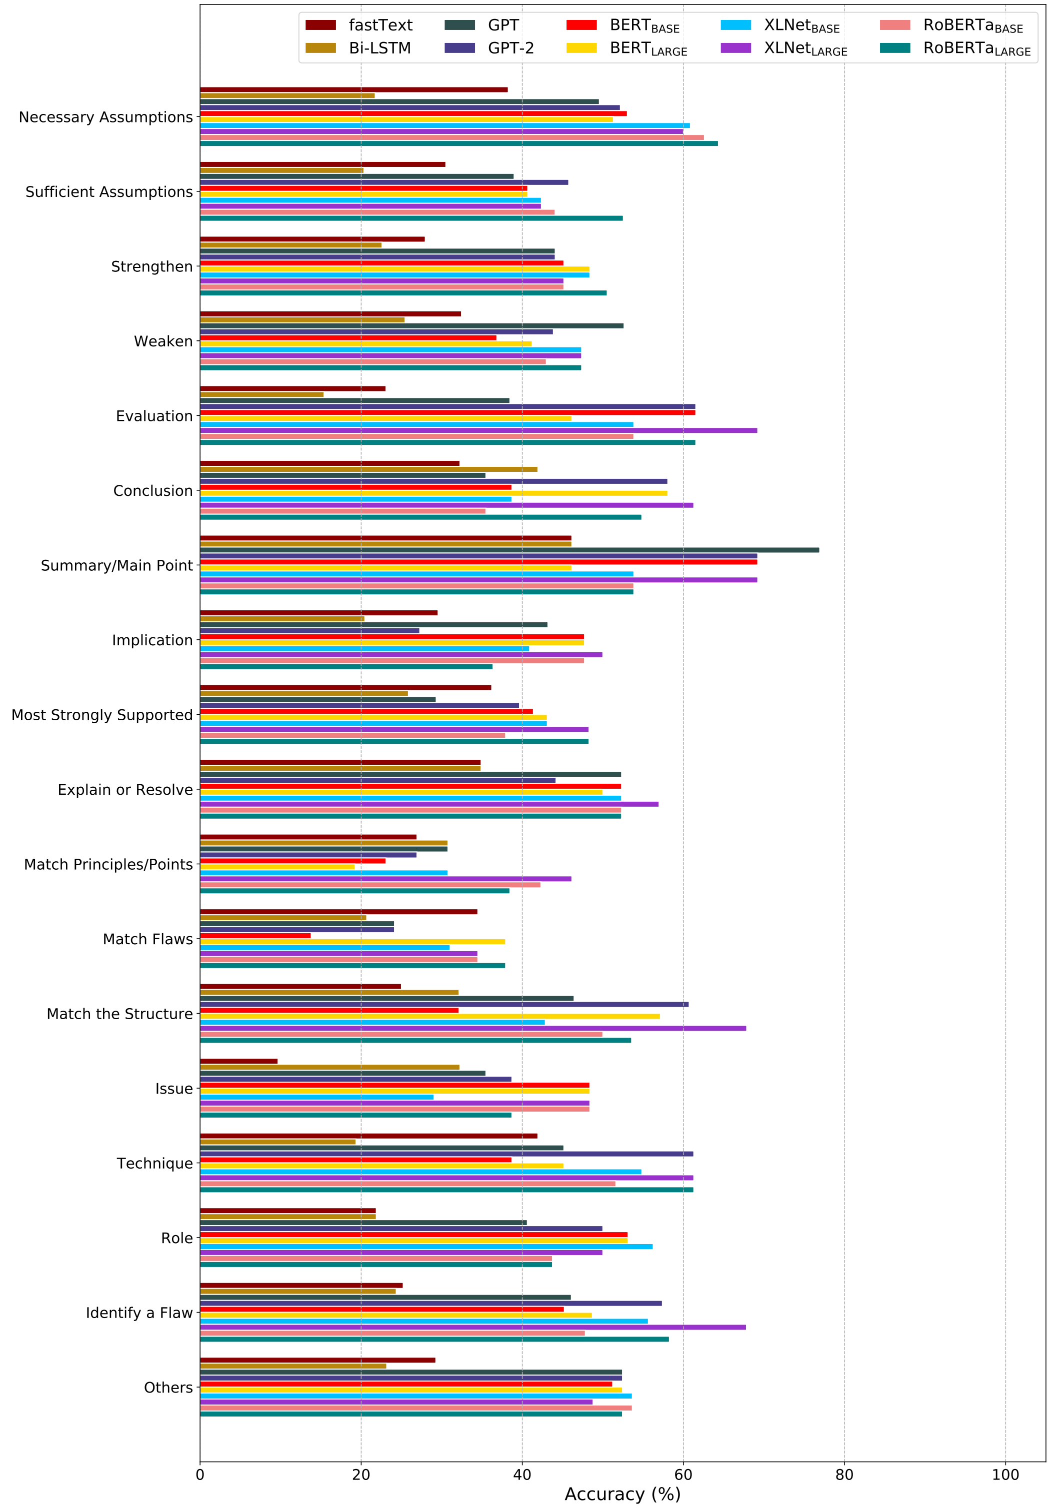


Figure 4: Accuracy of all baseline models on overall testing set

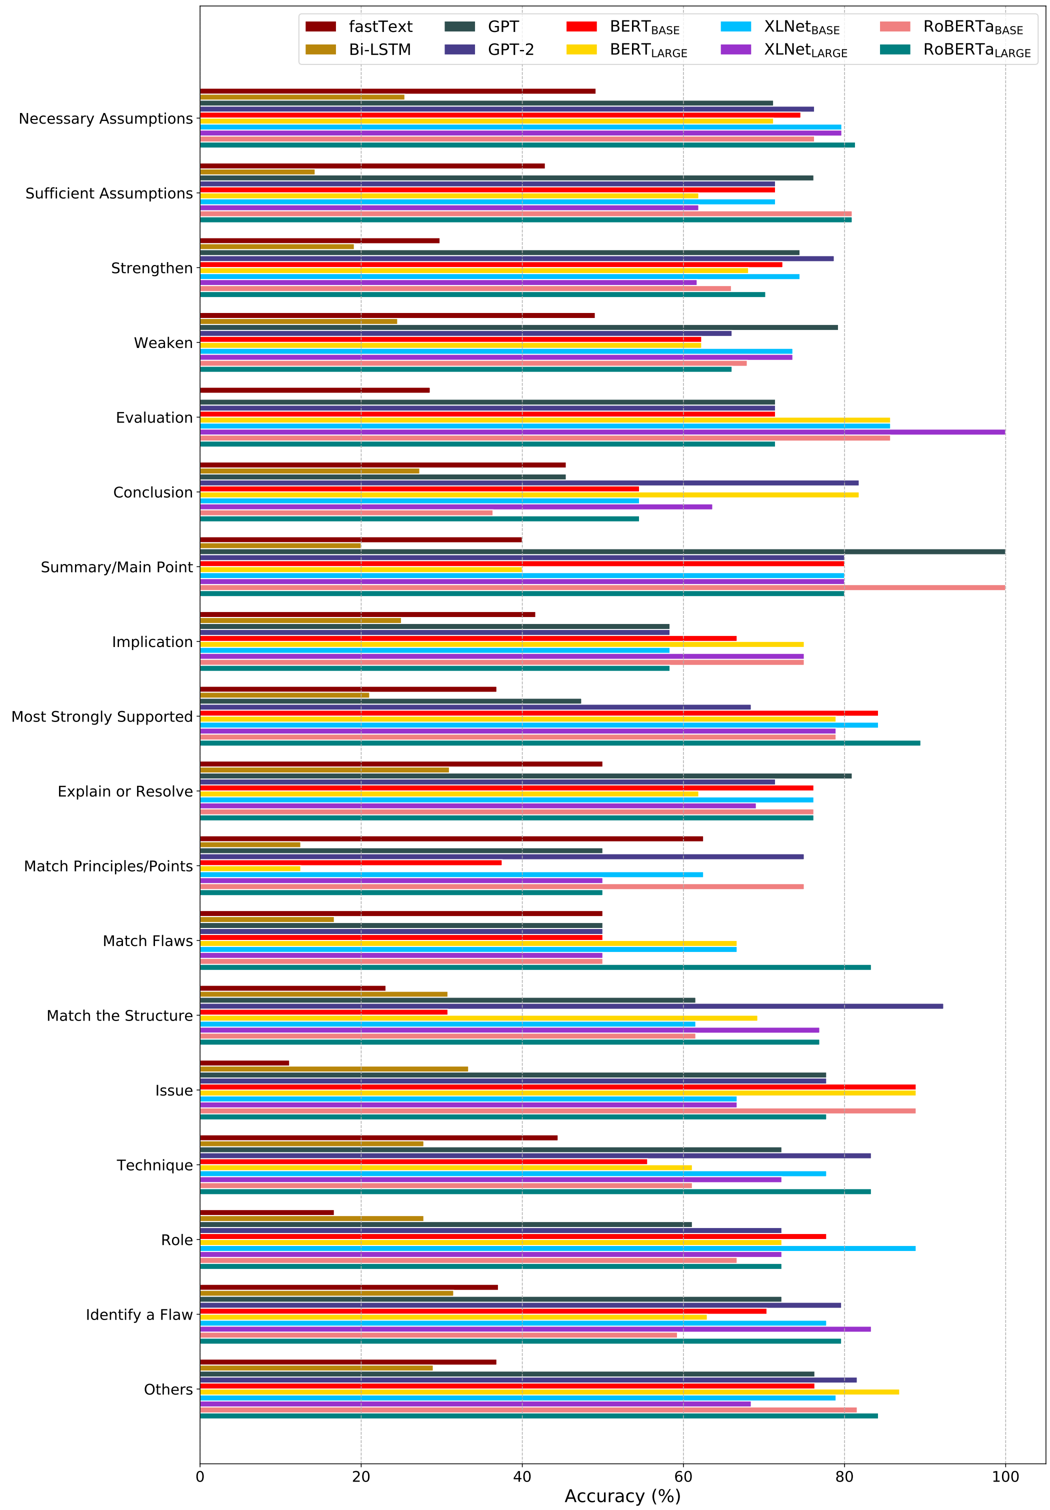


Figure 5: Accuracy of all baseline models on EASY set of testing set

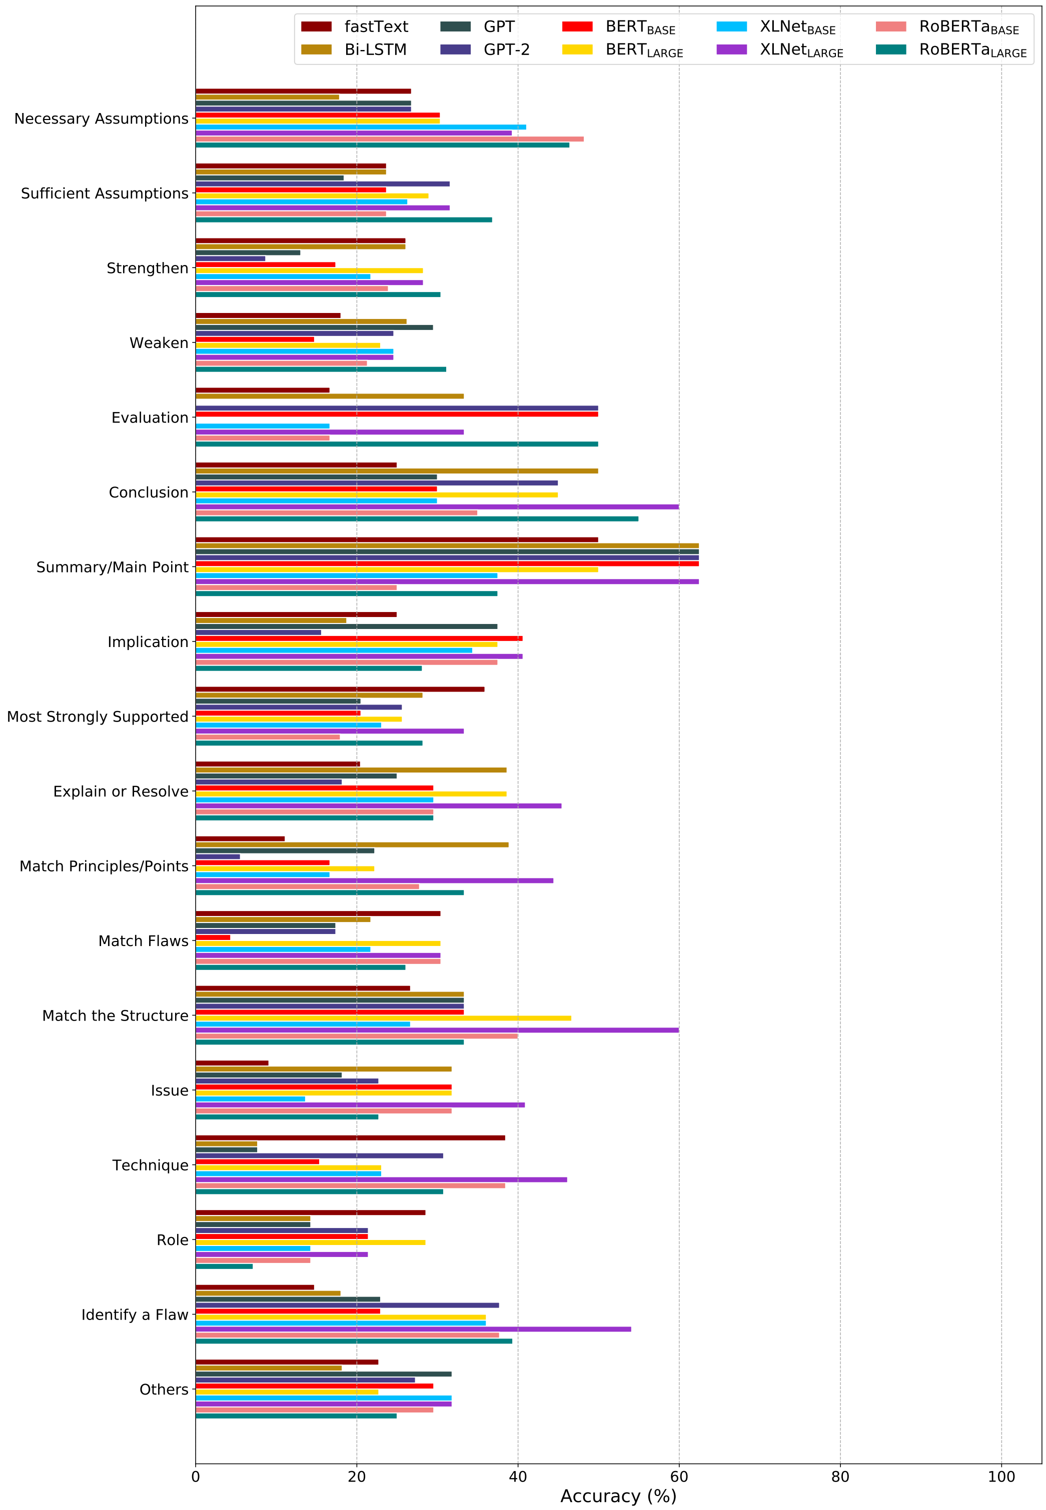


Figure 6: Accuracy of all baseline models on HARD set of testing set

| HYPERPARAM | GPT | GPT-2 | BERT$_{\text{BASE}}$ | BERT$_{\text{LARGE}}$ | XLNet$_{\text{BASE}}$ | XLNetT$_{\text{LARGE}}$ | RoBERTa$_{\text{BASE}}$ | RoBERTa$_{\text{LARGE}}$ |
|---|---|---|---|---|---|---|---|---|
| Learning Rate | 6.25e-5.0 | 6.25e-5 | 2e-05 | 1e-05 | 2e-05 | 5e-06 | 5e-06 | 5e-06 |
| Batch Size | 24 | 24 | 24 | 24 | 24 | 24 | 24 | 24 |
| Weight Decay | 0.01 | 0.01 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 |
| Learning Rate Decay | Linear | Linear | Linear | Linear | Linear | Linear | Linear | Linear |

Table 7: Hyperparameters for finetuning pre-training language models on ReClor

**Type:** Necessary Assumptions
**Definition:** identify the claim that must be true or is required in order for the argument to work

**Context:**
The current pattern of human consumption of resources, in which we rely on nonrenewable resources, for example metal ore, must eventually change. Since there is only so much metal ore available, ultimately we must either do without or turn to renewable resources to take its place.
**Question:**
Which one of the following is an assumption required by the argument?
**Options:**
A. We cannot indefinitely replace exhausted nonrenewable resources with other nonrenewable resources.
B. Consumption of nonrenewable resources will not continue to increase in the near future.
C. There are renewable resource replacements for all of the nonrenewable resources currently being consumed.
D. Ultimately we cannot do without nonrenewable resources.
**Answer:** A

Table 8: The definition and an example of the logical reasoning type - Necessary Assumptions

**Type:** Sufficient Assumptions
**Definition:** identify a sufficient assumption, that is, an assumption that, if added to the argument, would make it logically valid

**Context:**
Some theorists argue that literary critics should strive to be value-neutral in their literary criticism. These theorists maintain that by exposing the meaning of literary works without evaluating them, critics will enable readers to make their own judgments about the works' merits. But literary criticism cannot be completely value-neutral. Thus, some theorists are mistaken about what is an appropriate goal for literary criticism.
**Question:**
The argument's conclusion follows logically if which one of the following is assumed?
**Options:**
A. Any critic who is able to help readers make their own judgments about literary works' merits should strive to produce value-neutral criticism.
B. If it is impossible to produce completely value- neutral literary criticism, then critics should not even try to be value-neutral.
C. The less readers understand the meaning of a literary work, the less capable they will be of evaluating that work's merits.
D. Critics are more likely to provide criticisms of the works they like than to provide criticisms of the works they dislike.
**Answer:** B

Table 9: The definition and an example of the logical reasoning type - Sufficient Assumptions

| **Type:** Strengthen<br>**Definition:** identify information that would strengthen an argument |
|---|
| **Context:**<br>The television show Henry was not widely watched until it was scheduled for Tuesday evenings immediately after That' s Life, the most popular show on television. During the year after the move, Henry was consistently one of the ten most-watched shows on television. Since Henry' s recent move to Wednesday evenings, however, it has been watched by far fewer people. We must conclude that Henry was widely watched before the move to Wednesday evenings because it followed That' s Life and not because people especially liked it.<br>**Question:**<br>Which one of the following, if true, most strengthens the argument?<br>**Options:**<br>A. The show that now follows That's Life on Tuesdays has double the number of viewers it had before being moved.<br>B. Henry has been on the air for three years, but That's Life has been on the air for only two years.<br>C. After its recent move to Wednesday, Henry was aired at the same time as the second most popular show on television.<br>D. That's Life was not widely watched during the first year it was aired.<br>**Answer:** A |

Table 10: The definition and an example of the logical reasoning type - Strengthen

| **Type:** Weaken<br>**Definition:** identify information that would weaken an argument |
|---|
| **Context:**<br>When a certain gland becomes cancerous in humans, it produces high levels of a particular protein. A blood test can determine the level of this protein well before a cancer of the gland could be detected by other means. Some doctors recommend that aggressive anticancer treatment should be begun as early as possible for anyone who is tested and is found to have high levels of the protein.<br>**Question:**<br>Which one of the following, if true, most seriously weakens the doctors' recommendation?<br>**Options:**<br>A. The blood test for the protein has been in use for some time to monitor the condition of patients who have been diagnosed as having cancer of the gland.<br>B. Before the blood test became available, about one third of all cases of cancer of the gland were detected in early stages.<br>C. So far, no patients whose protein levels were found to be normal have subsequently developed cancer of the gland.<br>D. Enlargement of the gland, a common condition infrequently associated with cancer, results in high levels of the protein.<br>**Answer:** D |

Table 11: The definition and an example of the logical reasoning type - Weaken

| **Type:** Evaluation<br>**Definition:** identify information that would be useful to know to evaluate an argument |
|---|
| **Context:**<br>In a study, pairs of trained dogs were placed side by side and given a command such as šit. Äfter both obeyed the command, one dog was given a treat while its partner was given no reward at all. Over time, the dogs who went unrewarded began to disobey the command. This shows that dogs have an aversion to being treated unfairly.<br>**Question:**<br>Which one of the following would be most useful to know in order to evaluate the argument?<br>**Options:**<br>A. Were dogs who were accustomed to receiving regular rewards prior to the study more inclined to obey the command?<br>B. How many repetitions were required before the unrewarded dogs began to disobey the command?<br>C. Is there a decline in obedience if rewards are withheld from both dogs in the pair?<br>D. Were dogs who received treats in one trial ever used as dogs that did not receive treats in other trials?<br>**Answer:** C |

Table 12: The definition and an example of the logical reasoning type - Evaluation

| **Type:** Conclusion<br>**Definition:** identify the conclusion of a line of reasoning |
|---|
| **Context:**<br>Only an expert in some branch of psychology could understand why Patrick is behaving irrationally. But no expert is certain of being able to solve someone else' s problem. Patrick wants to devise a solution to his own behavioral problem.<br>**Question:**<br>Which one of the following conclusions can be validly drawn from the passage?<br>**Options:**<br>A. Patrick is not certain of being able to devise a solution to his own behavioral problem.<br>B. Unless Charles is an expert in some branch of psychology, Charles should not offer a solution to Patrick's behavioral problem.<br>C. If Charles is certain of being able to solve Patrick's behavioral problem, then Charles does not understand why Patrick is behaving in this way.<br>D. Patrick is not an expert in psychology.<br>**Answer:** C |

Table 13: The definition and an example of the logical reasoning type - Conclusion

| **Type:** Summary/Main Point<br>**Definition:** identify the main point of an argument |
|---|
| **Context:**<br>Balance is particularly important when reporting the 5. background of civil wars and conflicts. Facts must not be deliberately manipulated to show one party in a favorable light, and the views of each side should be fairly represented. This concept of balance, however, does not justify concealing or glossing over basic injustices in an effort to be even-handed. If all the media were to adopt such a perverse interpretation of balanced reporting, the public would be given a picture of a world where each party in every conflict had an equal measure of justice on its side, contrary to our experience of life and, indeed, our common sense.<br>**Question:**<br>Which one of the following best expresses the main point of the argument?<br>**Options:**<br>A. Balanced reporting presents the public with a picture of the world in which all sides to a conflict have equal justification.<br>B. Balanced reporting is an ideal that cannot be realized, because judgments of balance are necessarily subjective.<br>C. Common sense tells us that balance is especially needed when reporting the background of civil wars and conflicts.<br>D. Balanced reporting requires impartially revealing injustices where they occur no less than fairly presenting the views of each party in a conflict.<br>**Answer:** D |

Table 14: The definition and an example of the logical reasoning type - Summary/Main Point

| |
|---|
| **Type:** Implication<br>**Definition:** identify something that follows logically from a set of premises |
| **Context:**<br>The advanced technology of ski boots and bindings has brought a dramatic drop in the incidence of injuries that occur on the slopes of ski resorts: from 9 injuries per 1, 000 skiers in 1950 to 3 in 1980. As a result, the remainder of ski-related injuries, which includes all injuries occurring on the premises of a ski resort but not on the slopes, rose from 10 percent of all ski-related injuries in 1950 to 25 percent in 1980. The incidence of these injuries, including accidents such as falling down steps, increases with the amount of alcohol consumed per skier.<br>**Question:**<br>Which one of the following can be properly inferred from the passage?<br>**Options:**<br>A. Injuries that occurred on the slopes of ski resorts made up a smaller percentage of ski-related injuries in 1980 than in 1950.<br>B. As the number of ski injuries that occur on the slopes decreases, the number of injuries that occur on the premises of ski resorts increases.<br>C. If the technology of ski boots and bindings continues to advance, the incidence of ski-related injuries will continue to decline.<br>D. The technology of ski boots and bindings affects the incidence of each type of ski-related injury.<br>**Answer:** A |

Table 15: The definition and an example of the logical reasoning type - Implication

| |
|---|
| **Type:** Most Strongly Supported<br>**Definition:** find the choice that is most strongly supported by a stimulus |
| **Context:**<br>Light is registered in the retina when photons hit molecules of the pigment rhodopsin and change the molecules' shape. Even when they have not been struck by photons of light, rhodopsin molecules sometimes change shape because of normal molecular motion, thereby introducing error into the visual system. The amount of this molecular motion is directly proportional to the temperature of the retina.<br>**Question:**<br>Which one of the following conclusions is most strongly supported by the information above?<br>**Options:**<br>A. Molecules of rhodopsin are the only pigment molecules that occur naturally in the retina.<br>B. The visual systems of animals whose body temperature matches that of their surroundings are more error-prone in hot surroundings than in cold ones.<br>C. As the temperature of the retina rises, rhodopsin molecules react more slowly to being struck by photons.<br>D. The temperature of an animal's retina depends on the amount of light the retina is absorbing.<br>**Answer:** B |

Table 16: The definition and an example of the logical reasoning type - Most Strongly Supported

| |
|---|
| **Type:** Explain or Resolve<br>**Definition:**<br>identify information that would explain or resolve a situation<br>**Context:**<br>To reduce the mosquito population in a resort area, hundreds of trees were planted that bear fruit attractive to birds. Over the years, as the trees matured, they attracted a variety of bird species and greatly increased the summer bird population in the area. As expected, the birds ate many mosquitoes. However, the planting of the fruit trees had the very opposite of its intended effect.<br>**Question:**<br>Which one of the following, if true, most helps to explain the apparently paradoxical result?<br>**Options:**<br>A. Most of the species of birds that were attracted by the trees that were planted did not eat mosquitoes.<br>B. Increases and decreases in mosquito populations tend to follow a cyclical pattern.<br>C. The species of birds that were attracted in the greatest number by the fruit of the trees that were planted did not eat mosquitoes.<br>D. The birds attracted to the area by the trees ate many more insects that prey on mosquitoes than they did mosquitoes.<br>**Answer:** D |

Table 17: The definition and an example of the logical reasoning type - Explain or Resolve

| **Type:** Match Principles/Points<br>**Definition:** match principles from the stimulus to principles in a choice |
|---|
| **Context:**<br>Buying elaborate screensavers – programs that put moving images on a computer monitor to prevent damage – can cost a company far more in employee time than it saves in electricity and monitor protection. Employees cannot resist spending time playing with screensavers that flash interesting graphics across their screens.<br>**Question:**<br>Which one of the following most closely conforms to the principle illustrated above?<br>**Options:**<br>A. An electronic keyboard may be cheaper to buy than a piano but more expensive to repair.<br>B. An energy-efficient insulation system may cost more up front but will ultimately save money over the life of the house.<br>C. The time that it takes to have a pizza delivered may be longer than it takes to cook a complete dinner.<br>D. A complicated hotel security system may cost more in customer goodwill than it saves in losses by theft.<br>**Answer:** D |

Table 18: The definition and an example of the logical reasoning type - Match Principles/Points

| **Type:** Match Flaws<br>**Definition:** find a choice containing an argument that exhibits the same flaws as the passage's argument |
|---|
| **Context:**<br>The museum' s night security guard maintains that the thieves who stole the portrait did not enter the museum at any point at or above ground level. Therefore, the thieves must have gained access to the museum from below ground level.<br>**Question:**<br>The flawed pattern of reasoning in the argument above is most similar to that in which one of the following?<br>**Options:**<br>A. As had generally been expected, not all questionnaires were sent in by the official deadline. It follows that plans must have been made for the processing of questionnaires received late.<br>B. The store's competitors claim that the store, in selling off the shirts at those prices, neither made any profit nor broke even. Consequently, the store's customers must have been able to buy shirts there at less than the store's cost.<br>C. The product label establishes that this insecticide is safe for both humans and pets. Therefore, the insecticide must also be safe for such wild mammals as deer and rabbits.<br>D. If the census is to be believed, the percentage of men who are married is higher than the percentage of women who are married. Thus, the census must show a higher number of men than of women overall.<br>**Answer:** B |

Table 19: The definition and an example of the logical reasoning type - Match Flaws

| **Type:** Match the Structure<br>**Definition:** match the structure of an argument in a choice to the structure of the argument in the passage |
|---|
| **Context:**<br>It is an absurd idea that whatever artistic endeavor the government refuses to support it does not allow, as one can see by rephrasing the statement to read: No one is allowed to create art without a government subsidy.<br>**Question:**<br>The pattern of reasoning in which one of the following is most similar to that in the argument above?"<br>**Options:**<br>A. The notion that every scientist who has been supported by a government grant will be successful is absurd, as one can see by rewording it:No scientist is allowed to do research without a government grant.<br>B. The notion that every scientist who is supported by a government grant will be successful is absurd, as one can see by rewording it:No scientist lacking governmental support will be successful.<br>C. The claim that any driver who is not arrested does not break the law is absurd, as one can see by rewording it: Every driver who gets arrested has broken the law.<br>D. The claim that any driver who is not arrested does not break the law is absurd, as one can see by rewording it: Every driver who breaks the law gets arrested.<br>**Answer:** D |

Table 20: The definition and an example of the logical reasoning type - Match the Structure

| |
|---|
| **Type:** Issue<br>**Definition:** identify or infer an issue in dispute |
| **Context:**<br>Raphaela: Forcing people to help others is morally wrong. Therefore, no government has the right to redistribute resources via taxation. Anyone who wants can help others voluntarily. Edward: Governments do have that right, insofar as they give people the freedom to leave and hence not to live under their authority.<br>**Question:**<br>Raphaela and Edward disagree about the truth of which one of the following?<br>**Options:**<br>A. Any government that forces people to help others should permit emigration.<br>B. Any government that permits emigration has the right to redistribute resources via taxation.<br>C. Any government that redistributes resources via taxation forces people to help others.<br>D. Every government should allow people to help others voluntarily.<br>**Answer:** B |

Table 21: The definition and an example of the logical reasoning type - Issue

| |
|---|
| **Type:** Technique<br>**Definition:** identify the technique used in the reasoning of an argument |
| **Context:**<br>Joanna: The only way for a company to be successful, after emerging from bankruptcy, is to produce the same goods or services that it did before going bankrupt. It is futile for such a company to try to learn a whole new business. Ruth: Wrong. The Kelton Company was a major mining operation that went into bankruptcy. On emerging from bankruptcy, Kelton turned its mines into landfills and is presently a highly successful waste-management concern.<br>**Question:**<br>Ruth uses which one of the following argumentative techniques in countering Joanna's argument?<br>**Options:** A. She undermines a claim by showing that it rests on an ambiguity.<br>B. She offers an alternative explanation for a phenomenon.<br>C. She presents a counterexample to a claim.<br>D. She establishes a conclusion by excluding the only plausible alternative to that conclusion.<br>**Answer:** C |

Table 22: The definition and an example of the logical reasoning type - Technique

| |
|---|
| **Type:** Role<br>**Definition:** describe the individual role that a statement is playing in a larger argument |
| **Context:**The position that punishment should be proportional to how serious the offense is but that repeat offenders should receive harsher punishments than first-time offenders is unsustainable. It implies that considerations as remote as what an offender did years ago are relevant to the seriousness of an offense. If such remote considerations were relevant, almost every other consideration would be too. But this would make determining the seriousness of an offense so difficult that it would be impossible to apply the proportionality principle.<br><br>**Question:**<br>The statement that considerations as remote as what an offender did years ago are relevant to the seriousness of an offense plays which one of the following roles in the argument?<br>**Options:**<br>A. It is an allegedly untenable consequence of a view rejected in the argument's overall conclusion.<br>B. It is a statement the argument provides grounds to accept and from which the overall conclusion is inferred.<br>C. It is the overall conclusion in favor of which the argument offers evidence.<br>D. It is a premise offered in support of an intermediate conclusion of the argument.<br>**Answer:** A |

Table 23: The definition and an example of the logical reasoning type - Role

| |
|---|
| **Type:** Identify a Flaw<br>**Definition:** identify a flaw in an argument's reasoning |
| **Context:**<br>The tidal range at a particular location is the difference in height between high tide and low tide. Tidal studies have shown that one of the greatest tidal ranges in the world is found in the Bay of Fundy and reaches more than seventeen meters. Since the only forces involved in inducing the tides are the sun' s and moon' s gravity, the magnitudes of tidal ranges also must be explained entirely by gravitational forces.<br>**Question:**<br>Which one of the following most accurately describes a flaw in the reasoning above?<br>**Options:**<br>A. It does not differentiate between the tidal effect of the sun and the tidal effect of the moon.<br>B. It fails to consider that the size of a tidal range could be affected by the conditions in which gravitational forces act.<br>C. It presumes, without providing warrant, that most activity within the world's oceans is a result of an interplay of gravitational forces.<br>D. It gives only one example of a tidal range.<br>**Answer:** B |

Table 24: The definition and an example of the logical reasoning type - Identify a Flaw

| |
|---|
| **Type:** Others<br>**Definition:** other types of questions which are not included by the above |
| **Context:**<br>PhishCo runs a number of farms in the arid province of Nufa, depending largely on irrigation. Now, as part of a plan to efficiently increase the farms‘ total production, it plans to drill down to an aquifer containing warm, slightly salty water that will be used to raise fish in ponds. The water from the ponds will later be used to supplement piped-in irrigation water for PhishCo‘s vegetable fields, and the ponds and accompanying vegetation should help reduce the heat in the area of the farms.<br>**Question:**<br>Which of the following would, if true, most strongly suggest that the plan, if implemented, would increase the overall efficiency of PhishCo‘s farms?<br>**Options:**<br>A. Organic waste from fish in the pond water will help to fertilize fields where it is used for irrigation.<br>B. Fish raised on PhishCo's farms are likely to be saleable in the nearest urban areas.<br>C. Ponds will be located on low-lying land now partially occupied by grain crops.<br>D. The government of Nufa will help to arrange loan financing to partially cover the costs of drilling.<br>**Answer:** A |

Table 25: The definition and an example of the logical reasoning type - Others