

# When Choosing Plausible Alternatives, Clever Hans can be Clever

Pride Kavumba<sup>1,\*</sup>

Naoya Inoue<sup>1,2,\*</sup>

Benjamin Heinzerling<sup>2,1</sup>

Keshav Singh<sup>1</sup>

Paul Reisert<sup>2,1</sup>

Kentaro Inui<sup>1,2</sup>

<sup>1</sup>Tohoku University <sup>2</sup>RIKEN Center for Advanced Intelligence Project (AIP)  
{pkavumba, naoya-i, keshav.singh29, inui}@ecei.tohoku.ac.jp  
{benjamin.heinzerling, paul.reisert}@riken.jp

## Abstract

Pretrained language models, such as BERT and RoBERTa, have shown large improvements in the commonsense reasoning benchmark COPA. However, recent work found that many improvements in benchmarks of natural language understanding are not due to models learning the task, but due to their increasing ability to exploit superficial cues, such as tokens that occur more often in the correct answer than the wrong one. Are BERT’s and RoBERTa’s good performance on COPA also caused by this? We find superficial cues in COPA, as well as evidence that BERT exploits these cues. To remedy this problem, we introduce Balanced COPA, an extension of COPA that does not suffer from easy-to-exploit single token cues. We analyze BERT’s and RoBERTa’s performance on original and Balanced COPA, finding that BERT relies on superficial cues when they are present, but still achieves comparable performance once they are made ineffective, suggesting that BERT learns the task to a certain degree when forced to. In contrast, RoBERTa does not appear to rely on superficial cues.

## 1 Introduction

Pretrained language models such as ELMo (Peters et al., 2018), BERT (Devlin et al., 2019), and RoBERTa (Liu et al., 2019b) have led to improved performance in benchmarks of natural language understanding, in tasks such as natural language inference (NLI, Liu et al., 2019a), argumentation (Niven and Kao, 2019), and commonsense reasoning (Li et al., 2019; Sap et al., 2019). However, recent work has identified superficial cues in benchmark datasets which are predictive of the correct answer, such as token distributions and lexical overlap. Once these cues are neutralized, models perform poorly, suggesting that their good

\*Equal contribution.

The woman hummed to herself. What was the *cause* for this?

✓ She was in **a** good mood.

✗ She was nervous.

(a) Original COPA instance.

The woman trembled. What was the *cause* for this?

✗ She was in **a** good mood.

✓ She was nervous.

(b) Mirrored COPA instance.

Figure 1: A COPA instance (a) with premise and correct (✓) and wrong (✗) alternatives. Our analysis reveals that the unigram *a* (highlighted orange) is a superficial cue exploited by BERT. We neutralize such superficial cues by creating a mirrored instance (b). After mirroring, the highlighted superficial cue becomes ineffective in predicting the correct answer, since it occurs with equal probability in correct and wrong alternatives.

performance is an instance of the Clever Hans effect<sup>1</sup> (Pfungst, 1911): Models trained on datasets with superficial cues learn heuristics for exploiting these cues, but do not develop any deeper understanding of the task.

While superficial cues have been identified in, among others, datasets for NLI (Gururangan et al., 2018; McCoy et al., 2019), machine reading comprehension (Sugawara et al., 2018), and argumentation (Niven and Kao, 2019), one of the main benchmarks for commonsense reasoning, namely the Choice of Plausible Alternatives (COPA, Roemmele et al., 2011), has not been analyzed so far. Here we present an analysis of superficial cues in COPA.

<sup>1</sup>Named after the eponymous horse which appeared to be capable of simple mental tasks but actually relied on cues given involuntarily by its handler.

Given a premise, such as *The man broke his toe*, COPA requires choosing the more plausible, causally related alternative, in this case either: because *He got a hole in his sock* (wrong) or because *He dropped a hammer on his foot* (correct). To test whether COPA contains superficial cues, we conduct a dataset ablation in which we provide only partial input to the model. Specifically, we provide only the two alternatives, but not the premise, which makes solving the task impossible and hence should reduce the model to random performance. However, we observe that a model trained only on alternatives performs considerably better than random chance and trace this result to an unbalanced distribution of tokens between correct and wrong alternatives. Further analysis (§4.3) reveals that finetuned BERT (Devlin et al., 2019) perform very well (83.9 percent accuracy) on *easy* instances containing superficial cues, but worse (71.9 percent) on *hard* instances without such simple cues.

To prevent models from exploiting superficial cues in COPA, we introduce *Balanced COPA*. Balanced COPA contains one additional, *mirrored* instance for each original training instance. This mirrored instance uses the same alternatives as the corresponding original instance, but introduces a new premise which matches the *wrong* alternative of the original instance, e.g. *The man hid his feet*, for which the correct alternative is now because *He got a hole in his sock* (See another example in Figure 1). Since each alternative occurs exactly once as correct answer and exactly once as wrong answer in Balanced COPA, the lexical distribution between correct and wrong answers is perfectly balanced, i.e., superficial cues in the original alternatives have become uninformative.

Balanced COPA allows us to study the impact of the presence or absence of superficial cues on model performance.

Since BERT exploits cues in the original COPA, we expected performance to degrade when training on Balanced COPA. However, BERT trained on Balanced COPA performed comparably overall. As we will show, this is due to better performance on the “hard” instances. This suggests that once superficial cues are made uninformative, BERT learns the task to a certain degree.

In summary, our contributions are:

- We identify superficial cues in COPA that allow models to use simple heuristics instead

of learning the task (§2);

- We introduce Balanced COPA, which prevents models from exploiting these cues (§3);
- Comparing models on original and Balanced COPA, we find that BERT heavily exploits cues when they are present, but is also able to learn the task when they are not (§4); and
- We show that RoBERTa does not appear to exploit superficial cues.

## 2 Superficial Cues in COPA

### 2.1 COPA: Choice of Plausible Alternatives

Causal reasoning is an important prerequisite for natural language understanding. The Choice Of Plausible Alternatives (COPA) (Roemmele et al., 2011) is dataset that aims to benchmark causal reasoning in a simple binary classification setting.<sup>2</sup> COPA requires classifying sentence pairs consisting of the first sentence, the *premise*, and a second sentence that is either cause of, effect of, or unrelated to premise. Given the premise and two *alternatives*, one of which has a causal relation to the premise, while the other does not, models need to choose the more plausible alternative. Figure 1a shows an example of a COPA instance. The overall 1000 instances are split into training set<sup>3</sup> and test set of 500 instances each.

Prior to neural network approaches, the most dominant way of solving COPA was via Pointwise Mutual Information (PMI)-based statistics using a large background corpus between the content words in the premise and the alternatives (Gordon et al., 2011; Luo et al., 2016; Sasaki et al., 2017; Goodwin et al., 2012). Recent studies show that BERT and RoBERTa achieve considerable improvements on COPA (see Table 1).

However, recent work found that the strong performance of BERT and other deep neural models in benchmarks of natural language understanding can be partly or in some cases entirely explained by their capability to exploit superficial cues present in benchmark datasets. For example, Niven and Kao (2019) found that BERT exploits superficial cues, namely the occurrence of certain tokens such as *not*, in the Ar-

<sup>2</sup><http://people.ict.usc.edu/~gordon/copa.html>

<sup>3</sup>This set is called *development set* by Roemmele et al. (2011), but is used as training set by supervised models.

Model	Accuracy
BigramPMI (Goodwin et al., 2012)	63.4
PMI (Gordon et al., 2011)	65.4
PMI+Connectives (Luo et al., 2016)	70.2
PMI+Con.+Phrase (Sasaki et al., 2017)	71.4
BERT-large (Wang et al., 2019)	70.5
BERT-large (Sap et al., 2019)	75.0
BERT-large (Li et al., 2019)	75.4
RoBERTa-large (finetuned) <sup>4</sup>	90.6
BERT-large (finetuned)*	76.5 ± 2.7
RoBERTa-large (finetuned)*	87.7 ± 0.9

Table 1: Reported results on COPA. With the exception of (Wang et al., 2019), BERT-large and RoBERTa-large yields substantial improvements over prior approaches. See §2 for model details. \* indicates our replication experiments.

gment Reasoning Comprehension Task (Habernal et al., 2018). Similarly, Gururangan et al. (2018); Poliak et al. (2018); Dasgupta et al. (2018) showed that a simple text categorization model can perform well on the Stanford Natural Language Inference dataset (Bowman et al., 2015) and MultiNLI (Williams et al., 2018) when given incomplete input, even though the task should not be solvable without the full input. This suggests that the partial input contains unintended superficial cues that allow the models to take shortcuts without learning the actual task. Sugawara et al. (2018) investigated superficial cues that make questions easier across recent machine reading comprehension datasets. Given the fact that superficial cues were found in benchmark datasets for a wide variety of natural language understanding task, does COPA contain such cues, as well?

## 2.2 Token Distribution

One of the simplest types of superficial cues are unbalanced token distributions, i.e tokens appearing more often or less frequently with one particular instance label than with other labels. For example, Niven and Kao (2019) found that the token *not* occurs more often in one type of instance an argumentation dataset.

Similarly we identify superficial cues — in this case a single token that appears more frequently in correct alternatives or wrong alternatives — in the COPA training set. To find superficial cues in the form of predictive tokens, we use the following measures, defined by Niven and Kao (2019). Let  $\mathbb{T}_j^{(i)}$  be the set of tokens in the alternatives for

data point  $i$  with label  $j$ . The *applicability*  $\alpha_k$  of a token  $k$  counts how often this token occurs in an alternative with one label, but not the other:

$$\alpha_k = \sum_{i=1}^n \mathbb{1} \left[ \exists j, k \in \mathbb{T}_j^{(i)} \wedge k \notin \mathbb{T}_{-j}^{(i)} \right]$$

The *productivity*  $\pi_k$  of a token is the proportion of applicable instances for which it predicts the correct answer:

$$\pi_k = \frac{\sum_{i=1}^n \mathbb{1} \left[ \exists j, k \in \mathbb{T}_j^{(i)} \wedge k \notin \mathbb{T}_{-j}^{(i)} \wedge y_i = j \right]}{\alpha_k}$$

Finally, the *coverage*  $\xi_k$  of a token is the proportion of applicable instances among all instances:

$$\xi_k = \frac{\alpha_k}{n}$$

Table 2 shows the five tokens with highest coverage. For example, *a* is the token with the highest coverage and appears in either a correct alternative or wrong alternative in 21.2% of COPA training instances. Its productivity of 57.5% expresses that it appears in in correct alternatives 7.5% more often than expected by random chance. This suggests that a model could rely on such unbalanced distributions of tokens to predict answers based only on alternatives without understanding the task.

To test this hypothesis, we perform a dataset ablation, providing only the two alternatives as input to RoBERTa, but not the premise, following similar ablations by Gururangan et al. (2018); Niven and Kao (2019). RoBERTa trained<sup>5</sup> in this setting, i.e. on alternatives only, achieves a mean accuracy of 59.6 (± 2.3). This is problematic because COPA is designed as a choice between alternatives given the premise. Without a premise given, model performance should not exceed random chance. Consequently, a result better than random chance shows that the dataset allows solving the task in a way that was not intended by its creators. To fix this problem, we create a balanced version of COPA that does not suffer from unbalanced token distributions in correct and wrong alternatives.

## 3 Balanced COPA (B-COPA)

To allow evaluating models on a benchmark without superficial cues, we need to make them inef-

<sup>5</sup>See §4.1 for experimental setup.

Cue	App.	Prod.	Cov.
in	47	55.3	9.40
was	55	61.8	11.0
to	82	40.2	16.4
the	85	38.8	17.0
a	106	57.5	21.2

Table 2: Applicability (App.), Productivity (Prod.) and Coverage (Cov.) of the various words in the *alternatives* of the COPA dev set.

fective. Our approach is to balance the token distributions in correct alternatives and wrong alternatives in the training set. Without unbalanced token distributions, we hope models are able to learn other patterns more closely related to the task, e.g. a pair of causally related events, rather than superficial cues.

### 3.1 Data Collection

To create the balanced COPA training set, we manually mirror the original training set by modifying the premise. Taking the original training set as a starting point, we duplicate the COPA instances and modify their premises so that incorrect alternatives become correct. Suppose the following original COPA instance:

- **Premise:** The stain came out of the shirt. What was the CAUSE of this?
- **Alternative 1:** I bleached the shirt. (Correct)
- **Alternative 2:** I patched the shirt.

We create the following balanced COPA instance, where the wrong alternative becomes the correct choice now:

- **Premise:** *The shirt did not have a hole anymore.* What was the CAUSE of this?
- **Alternative 1:** I bleached the shirt.
- **Alternative 2:** I patched the shirt. (Correct)

This approach is similar to [Niven and Kao \(2019\)](#), who create a balanced benchmark of the Argument Reasoning Comprehension Task by negating and rotating its ingredients, exploiting the nature of the task. However, due to the nature of COPA, we cannot follow their approach and choose to create new premises.

Dataset	Accuracy	Fleiss’ kappa $k$
Original COPA	100.0	0.973
Balanced COPA	97.0	0.798

Table 3: Results of human performance evaluation of the original COPA and Balanced COPA.

To collect such balanced data, we asked five fluent English speakers who have background knowledge of NLP (see [Appendix A](#) for the detailed guideline). Finally, we collected 500 new mirrored instances. Concatenating it with the original training instances, the balanced COPA consists of 1,000 instances in total. The corpus is publicly available at <https://balanced-copa.github.io>.

### 3.2 Quality Evaluation

To ensure the quality of the mirrored instances, we estimate a human performance using Amazon Mechanical Turk (AMT), a widely-used crowdsourcing platform. We randomly sample 100 instances from the original COPA training set and 100 instances from the balanced COPA, and asked crowdworkers to solve each instance (see [Appendix B](#) for an actual screenshot). To avoid noisy workers, we presented our tasks to workers who meet master AMT qualification with at least 10,000 HIT approvals and 99% HIT approval rate. Per HIT, we assign three crowd workers and offer 10 cents reward.

From the collected responses, we calculate the accuracy of workers (by majority voting) and inter-annotator agreement by Fleiss’ Kappa ([Fleiss, 1981](#)). The human evaluation shows that our mirrored instances are comparable in difficulty to the original ones (see [Table 3](#)). However, we found that some mirrored instances are a bit tricky at first glance. But, with a bit more attention, the answer is quite obvious (see [Appendix C](#), for an example).

## 4 Experiments

### 4.1 BERT and RoBERTa on COPA

In this section we analyze the performance of two recent pretrained language models on COPA: BERT and RoBERTa, an optimized variant of BERT that achieves better performance on the SuperGLUE benchmark ([Wang et al., 2019](#)), which includes COPA.

Model	Method	Training Data	Overall	Easy	Hard	p-value (%)
Goodwin et al. (2012)	PMI	unsupervised	61.8	64.7	60.0	19.8
Gordon et al. (2011)	PMI	unsupervised	65.4	65.8	65.2	83.5
Sasaki et al. (2017)	PMI	unsupervised	71.4	75.3	69.0	4.8*
Word frequency	wordfreq	COPA	53.5	57.4	51.3	9.8
BERT-large-FT	LM, NSP	COPA	76.5 ( $\pm 2.7$ )	83.9 ( $\pm 4.4$ )	71.9 ( $\pm 2.5$ )	0.0*
RoBERTa-large-FT	LM	COPA	87.7 ( $\pm 0.9$ )	91.6 ( $\pm 1.1$ )	85.3 ( $\pm 2.0$ )	0.0*

Table 4: Model performance on the COPA test set (*Overall*), on *Easy* instances with superficial cues, and on *Hard* instances without superficial cues. p-values according to Approximate Randomization Tests (Noreen, 1989), with \* indicating a significant difference between performance on *Easy* and *Hard*  $p < 5\%$ . Methods are pointwise mutual information (PMI), word frequency provided by the wordfreq package (Speer et al., 2018), pretrained language model (LM), and next-sentence prediction (NSP).

We convert COPA instances as follows to make them compatible with the input format required by BERT/RoBERTa. For a COPA instance  $\langle p, a_1, a_2, q \rangle$ , where  $p$  is a premise,  $a_i$  is the  $i$ -th alternative, and  $q$  is a question type (either *effect* or *cause*), we construct BERT’s input depending on the question type. We assume that the first sentence and the second sentence in the next sentence prediction task describe a cause and an effect, respectively. Specifically, for each  $i$ -th alternative, we define the following input function:

$$\text{input}(p, a_i) = \begin{cases} \text{“[CLS] } p \text{ [SEP] } a_i \text{ [SEP]”} & \text{if } q \text{ is effect} \\ \text{“[CLS] } a_i \text{ [SEP] } p \text{ [SEP]”} & \text{if } q \text{ is cause} \end{cases}$$

Part of BERT’s training objective includes next sentence prediction. Given a pair of sentences, BERT predicts whether one sentence can be plausibly followed by the other. For this, BERT’s input format contains two [SEP] tokens to mark the two sentences and the [CLS] token, which is used as the input representation for next sentence prediction. This part of BERT’s architecture makes it a natural fit for COPA.

One of the key differences between BERT and RoBERTa is that the next sentence prediction objective is not part of RoBERTa’s training objective. Instead, RoBERTa is trained with masked language modeling only, with its input consisting of multiple concatenated sentences. To match this training setting, we encode two sentences in a single segment as follows:

$$\text{input}(p, a_i) = \begin{cases} \text{“<s> } p \text{ } a_i \text{ </s>}” & \text{if } q \text{ is effect} \\ \text{“<s> } a_i \text{ } p \text{ </s>}” & \text{if } q \text{ is cause} \end{cases}$$

After encoding premise-alternative with BERT or RoBERTa, we take the first hidden representation  $z_i^0$ , i.e. the one corresponding to [CLS] or

<s>, in the final model layer and pass it through a linear layer for binary classification:

$$y_i = \mathbf{w}^\top \mathbf{z}_i^0 + b, \quad (1)$$

where the parameters  $\mathbf{w} \in \mathbb{R}^h$  and  $b \in \mathbb{R}$  are learned on the COPA training set. Finally, we choose the alternative with the higher score, i.e.,  $a_{\hat{i}}$  with  $\hat{i} = \arg \max_{i \in \{1,2\}} y_i$ .

For training, we minimize the cross entropy loss with the logits  $[y_1; y_2]$  and fine-tune BERT and RoBERTa’s parameters. In our experiments, we use pretrained BERT-large (uncased) with 24 layers, 16 self-attention heads (340M parameters) and pretrained RoBERTa-large with 24 layers, 16 self-attention heads (355M parameters).<sup>6</sup>

## 4.2 Training Details

For training, we consider two configurations: (i) using the original COPA training set (§4.3), and (ii) using B-COPA (§4.4). We randomly split the training data into training data and validation data with the ratio of 9:1. For B-COPA, we make sure that a pair of original instance and its mirrored counterpart always belong to the same split in order to ensure that a model is trained without superficial cues. For testing, we use all 500 instances from the original COPA test set.

We run each experiment three times with different random seeds and average the results. We train for 10 epochs and choose the best model based on the validation score. To reduce GPU RAM usage, we set BERT and RoBERTa’s maximum sequence length to 32, which covers all training and test instances. We use Adam (Kingma and Ba, 2015) with warmup, weight decay of 0.01, a batch size

<sup>6</sup><https://huggingface.co/pytorch-transformers/>

Model	Training data	Overall	Easy	Hard
BERT-large-FT	B-COPA	74.5 ( $\pm$ 0.7)	74.7 ( $\pm$ 0.4)	<b>74.4</b> ( $\pm$ 0.9)
BERT-large-FT	B-COPA (50%)	74.3 ( $\pm$ 2.2)	76.8 ( $\pm$ 1.9)	72.8 ( $\pm$ 3.1)
BERT-large-FT	COPA	<b>76.5</b> ( $\pm$ 2.7)	<b>83.9</b> ( $\pm$ 4.4)	71.9 ( $\pm$ 2.5)
RoBERTa-large-FT	B-COPA	<b>89.0</b> ( $\pm$ 0.3)	88.9 ( $\pm$ 2.1)	<b>89.0</b> ( $\pm$ 0.8)
RoBERTa-large-FT	B-COPA (50%)	86.1 ( $\pm$ 2.2)	87.4 ( $\pm$ 1.1)	85.4 ( $\pm$ 2.9)
RoBERTa-large-FT	COPA	87.7 ( $\pm$ 0.9)	<b>91.6</b> ( $\pm$ 1.1)	85.3 ( $\pm$ 2.0)

Table 5: Results of fine-tuned models on Balanced COPA. *Easy*: instances with superficial cues, *Hard*: instances without superficial cues.

of 4, and a gradient accumulation of 8. We optimize hyperparameters for BERT and RoBERTa separately on the validation set. For BERT, we test learning rates of  $2e-4$ ,  $1e-4$ ,  $8e-5$ ,  $4e-5$ ,  $2e-5$ , and  $1e-5$ , and use warm up proportion of 0.1, with gradient norm clipping of 1.0. For RoBERTa, we test learning rates of  $1e-5$ ,  $8e-6$ ,  $6e-6$ ,  $4e-6$ ,  $2e-6$ , and  $1e-6$ , and use warm up proportion of 0.06, with no gradient norm clipping.

### 4.3 Evaluation on Easy and Hard subsets

To investigate the behaviour of BERT and RoBERTa trained on the original COPA, which contains superficial cues, we split the test set into an *Easy* subset and a *Hard* subset. The *Easy subset* consists of instances that are correctly solved by the premise-oblivious model described in §2. To account for variation between the three runs with different random seeds, we deem an instance correctly classified only if the premise-oblivious model’s prediction is correct for all three runs. This results in the *Easy* subset with 190 instances and the *Hard* subset comprising the remaining 310 instances. Such an easy/hard split follows similar splits in NLI datasets (Gururangan et al., 2018).

We then compare BERT and RoBERTa with previous models on the *Easy* and *Hard* subsets.<sup>7</sup> As Table 4 shows, previous models perform similarly on both subsets, with the exception of Sasaki et al. (2017).<sup>8</sup> Overall both BERT (76.5%) and RoBERTa (87.7%) considerably outperform the

<sup>7</sup>For previous models, we use the prediction keys available on <http://people.ict.usc.edu/~gordon/copa.html>

<sup>8</sup>We conjecture that word frequency is another superficial cue exploited by models. To verify this we train a classifier based on word frequencies only (Speer et al., 2018) and find that this classifier is able to identify the correct alternative better than random chance, but this result is not significant ( $p = 9.8\%$ ).

best previous model (71.4%). However, BERT’s improvements over previous work can be almost entirely attributed to high accuracy on the *Easy* subset: on this subset, finetuned BERT-large improves 8.6 percent over the model by (Sasaki et al., 2017) (83.9% vs. 75.3%), but on the *Hard* subset, the improvement is only 2.9 percent (71.9% vs. 69.0%). This indicates that BERT relies on superficial cues. The difference between accuracy on *Easy* and *Hard* is less pronounced for RoBERTa, but still suggests some reliance on superficial cues. We speculate that superficial cues in the COPA training set prevented BERT and RoBERTa from focusing on task-related non-superficial cues such as causally related event pairs.

### 4.4 Evaluation on Balanced COPA (B-COPA)

How will BERT and RoBERTa behave when there are no superficial cues in the training set? To answer this question, we now train BERT and RoBERTa on B-COPA and evaluate on the *Easy* and *Hard* subsets. The results are shown in Table 5. The smaller performance gap between *Easy* and *Hard* subsets indicates that training on B-COPA encourages BERT and RoBERTa to rely less on superficial cues. Moreover, training on B-COPA improves performance on the *Hard* subset, both when training with all 1000 instances in B-COPA, and when matching the training size of the original COPA (500 instances, *B-COPA 50%*). Note that training on *B-COPA 50%* exposes the model to lexically less diverse training instances than the original COPA due to the high overlap between mirrored alternatives (see §3).

These results show that once superficial cues are removed, the models are able to learn the task to a high degree. This contrasts with Niven and Kao (2019), who found that BERT’s performance

Model	Training data	Overall	Easy	Hard
BERT-large	B-COPA	70.5 ( $\pm$ 2.5)	72.6 ( $\pm$ 2.3)	<b>69.1</b> ( $\pm$ 2.7)
BERT-large	B-COPA (50%)	69.9 ( $\pm$ 1.9)	71.2 ( $\pm$ 1.3)	69.0 ( $\pm$ 3.5)
BERT-large	COPA	<b>71.7</b> ( $\pm$ 0.5)	<b>80.5</b> ( $\pm$ 0.4)	66.3 ( $\pm$ 0.8)
RoBERTa-large	B-COPA	<b>76.7</b> ( $\pm$ 0.8)	73.3 ( $\pm$ 1.5)	<b>78.8</b> ( $\pm$ 2.0)
RoBERTa-large	B-COPA (50%)	72.4 ( $\pm$ 2.0)	72.1 ( $\pm$ 1.7)	72.6 ( $\pm$ 2.1)
RoBERTa-large	COPA	76.4 ( $\pm$ 0.7)	<b>79.6</b> ( $\pm$ 1.0)	74.4 ( $\pm$ 1.1)
BERT-base-NSP	None	<b>66.4</b>	66.2	<b>66.7</b>
BERT-large-NSP	None	65.0	<b>66.9</b>	62.1

Table 6: Results of non-fine-tuned models on Balanced COPA. *Easy*: instances with superficial cues, *Hard*: instances without superficial cues.

on the Argument Reasoning Comprehension Task (Habernal et al., 2018) does not exceed random chance level after superficial cues are made uninformative. A likely explanation for this contrast is the difference in the inherent task difficulties. Argument reasoning comprehension is a high level natural language understanding task requiring world knowledge and complex reasoning skills, while COPA can be largely solved with associative reasoning, as the performance of the PMI-based baselines shows (Table 4). A second possible explanation is BERT’s insensitivity to negations (Ettinger, 2019). Since Niven and Kao (2019) made superficial cues uninformative by adding negated instances to the dataset, BERT’s insensitivity to negations makes distinguishing between instances and negated instances difficult (see §3).

#### 4.5 Analysis of sentence pair embeddings

The findings presented in the previous sections, namely BERT’s and RoBERTa’s good performance on COPA in spite of the rather small amount of training data, leads us to the following hypothesis that pretraining enables these models to create an embedding space in which embeddings of plausible sentence pairs are distinguishable from embeddings of less plausible pairs.

To investigate how well the respective embedding spaces of BERT and RoBERTa separate plausible and less-plausible pairs, we train BERT-large and RoBERTa-large *without fine-tuning*. Specifically, we freeze model weights and train a classifier by parameterizing  $w$  and  $b$  in Equation 1 as a soft-margin Support Vector Machine (SVM,

Cortes and Vapnik, 1995).<sup>9</sup> We also report results for a simple model that only uses BERT’s pre-trained next sentence predictor (BERT-base-NSP, BERT-large-NSP), i.e., we choose the alternative with the higher next sentence prediction score. The results are shown in Table 6. The relatively high accuracies of BERT-large, RoBERTa-large and BERT\*-NSP show that these pretrained models are already well-equipped to perform this task “out-of-the-box”.

#### 4.6 Analysis of sensitivity to cues

To analyze the sensitivity of BERT and RoBERTa to superficial cues and to content words, we employ a gradient-based approach, following (Bruner et al., 2019). Specifically, we define the sensitivity  $s_{i,t}$  of the classification score in  $i$ -th COPA test instance to input token  $t$ , as follows:

$$s_{i,t} = \frac{\|\mathbf{g}_t\|}{\sum_{t' \in T_i} \|\mathbf{g}_{t'}\|}, \mathbf{g}_t = \frac{\partial y}{\partial \mathbf{x}_t}, \quad (2)$$

where  $T_i$  is a sequence of all input tokens in the  $i$ -th COPA test instance,  $y$  is a score function defined by Equation (1), and  $\mathbf{x}_t \in \mathbb{R}^{1024}$  is a position-augmented token embedding of  $t$ . We then define the sensitivity  $S(k)$  to cue  $k$  over all COPA test instances as the average over all  $m$  COPA test instances:  $S(k) = \frac{1}{m} \sum_i^m s_{i,k}$ .

We are interested in the change of sensitivity towards cue  $t$  of a model trained on original COPA compared to a model trained on Balanced COPA. We plot this difference as a function of the cue’s productivity (Figure 2). We observe that BERT trained on Balanced COPA is less sensitive to a

<sup>9</sup>We tune the SVM hyperparameter  $C \in \{0.0001, 0.001, 0.01, 0.1, 1\}$  on the validation set.

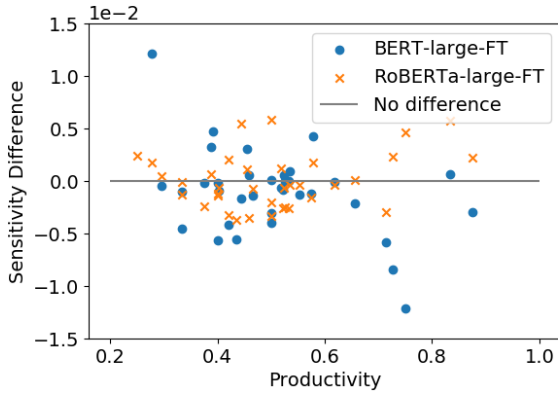


Figure 2: Change of sensitivity to superficial cues (in §2) from COPA-trained models to B-COPA-trained models as a function of their productivity.

Cue	$S^{\text{COPA}}$	$S^{\text{B-COPA}}$	Diff.	Prod.
woman	7.98	4.84	-3.14	0.25
mother	5.16	3.95	-1.21	0.75
went	6.00	5.15	-0.85	0.73
down	5.52	4.93	-0.58	0.71
into	4.07	3.51	-0.56	0.40

Table 7: Sensitivity of BERT-large to superficial cues identified in §2 (unit:  $10^{-2}$ ). Cues with top-5 reduction are shown.  $S^{\text{COPA}}$ ,  $S^{\text{B-COPA}}$  indicate the mean contributions of BERT-large trained on COPA, and BERT-large trained on B-COPA, respectively.

few highly productive superficial cues than BERT trained on original COPA. Note the decrease in the sensitivity for cues of productivity from 0.7 to 0.9. These cues are shown in Table 7. However, for cues with lower productivity, the picture is less clear, in case of RoBERTa, there are no noticeable trends in the change of sensitivity.

## 5 Conclusions

We established that COPA, an important benchmark of commonsense reasoning, contains superficial cues, specifically single tokens predictive of the correct answer, that allow models to solve the task without actually understanding it. Our experiments suggest that BERT’s good performance on COPA can be explained by its ability to exploit these superficial cues. BERT performs well on *Easy* instances with such superficial cues, and comparable to previous methods on *Hard* instances without such cues. RoBERTa, in contrast, represents a real improvement considerably outperforms both BERT and previous methods on *Hard* instances as well.

To allow evaluating models on a benchmark without predictive single tokens, we created the Balanced COPA dataset. Balanced COPA neutralizes this kind of superficial cue by mirroring instances from the original COPA dataset, thereby removing any differences in token distributions between correct and wrong alternatives. Surprisingly, we found that both BERT and RoBERTa finetuned on Balanced COPA perform comparably overall to the models finetuned on the original COPA. However, a more detailed analysis revealed quite different behaviour. Whereas BERT finetuned on original COPA heavily exploited superficial cues, we now find evidence that BERT finetuned on balanced COPA appears to learn some aspects of the task with similar accuracies on both *Easy* and *Hard* instances. Even more surprisingly, RoBERTa benefits from training on Balanced COPA instances and achieves higher accuracy than on the original COPA with superficial cues.

Two important questions remain unanswered at present, which we plan to explore in future work: Even in the presence of superficial cues, RoBERTa does not seem to rely on them. First, why does RoBERTa not appear to rely on superficial cues, even when they are available? And second, are the results of our experiments on Balanced COPA specific to BERT and RoBERTa or are all pretrained language models able to exploit superficial cues in COPA and able to solve the task by other means if no such cues are present?

## Acknowledgements

This work was partially supported by JSPS KAKENHI Grant Number 19K20332 and JST CREST Grant Number JPMJCR1513, including AIP challenge.

## References

- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics.
- Gino Brunner, Yang Liu, Dami’an Pascual, Oliver Richter, and Roger Wattenhofer. 2019. On the validity of self-attention as explanation in transformer models. *ArXiv*, abs/1908.04211.



- Corinna Cortes and Vladimir Vapnik. 1995. [Support-vector networks](#). *Mach. Learn.*, 20(3):273–297.
- Ishita Dasgupta, Demi Guo, Andreas Stuhlmüller, Samuel J Gershman, and Noah D. Goodman. 2018. Evaluating compositionality in sentence embeddings. *ArXiv*, abs/1802.04302.
- 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.
- Allyson Ettinger. 2019. [What bert is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#).
- Joseph L. Fleiss. 1981. *Statistical methods for rates and proportions*, 2nd edition. Wiley, New York.
- Travis Goodwin, Bryan Rink, Kirk Roberts, and Sanda Harabagiu. 2012. [UTDHLT: COPACETIC system for choosing plausible alternatives](#). In *\*SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 461–466, Montréal, Canada. Association for Computational Linguistics.
- Andrew S. Gordon, Cosmin Adrian Bejan, and Kenji Sagae. 2011. [Commonsense Causal Reasoning Using Millions of Personal Stories](#). In *25th Conference on Artificial Intelligence (AAAI-11)*, San Francisco, CA.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel Bowman, and Noah A. Smith. 2018. [Annotation artifacts in natural language inference data](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 107–112, New Orleans, Louisiana. Association for Computational Linguistics.
- Ivan Habernal, Henning Wachsmuth, Iryna Gurevych, and Benno Stein. 2018. [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)*, pages 1930–1940, New Orleans, Louisiana. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Zhongyang Li, Tongfei Chen, and Benjamin Van Durme. 2019. [Learning to rank for plausible plausibility](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4818–4823, Florence, Italy. Association for Computational Linguistics.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019a. Multi-task deep neural networks for natural language understanding. *arXiv preprint arXiv:1901.11504*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019b. [Roberta: A robustly optimized BERT pretraining approach](#). *CoRR*, abs/1907.11692.
- Zhiyi Luo, Yuchen Sha, Kenny Q. Zhu, Seung-won Hwang, and Zhongyuan Wang. 2016. [Commonsense causal reasoning between short texts](#). In *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning, KR’16*, pages 421–430. AAAI Press.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Timothy Niven and Hung-Yu Kao. 2019. [Probing neural network comprehension of natural language arguments](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Eric W Noreen. 1989. *Computer-intensive methods for testing hypotheses*. Wiley New York.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Oskar Pfungst. 1911. *Clever Hans:(the horse of Mr. Von Osten.) a contribution to experimental animal and human psychology*. Holt, Rinehart and Winston.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis only baselines in natural language inference](#). In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.

Melissa Roemmele, Cosmin Adrian Bejan, and Andrew S Gordon. 2011. [Choice of plausible alternatives: An evaluation of commonsense causal reasoning](#). In *AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning*, Stanford University.

Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. Socialiqa: Commonsense reasoning about social interactions. *ArXiv*, abs/1904.09728.

Shota Sasaki, Sho Takase, Naoya Inoue, Naoaki Okazaki, and Kentaro Inui. 2017. [Handling multiword expressions in causality estimation](#). In *IWCS 2017 — 12th International Conference on Computational Semantics — Short papers*.

Robyn Speer, Joshua Chin, Andrew Lin, Sara Jewett, and Lance Nathan. 2018. [Luminosoin-sight/wordfreq: v2.2](#).

Saku Sugawara, Kentaro Inui, Satoshi Sekine, and Akiko Aizawa. 2018. [What makes reading comprehension questions easier?](#) In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4208–4219, Brussels, Belgium. Association for Computational Linguistics.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. SuperGLUE: A stickier benchmark for general-purpose language understanding systems. *arXiv preprint 1905.00537*.

Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. [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)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.

## A Balanced COPA New Premise Guidelines

We instructed dataset creators with the following guidelines:

1. Ensure as much lexical overlap in new premise as the original premise.
2. Ensure little lexical overlap between premise and alternative, but if a word occurs both in premise and alternatives, it is acceptable to include it in the premise.
3. Maintain, as much as possible, the length and style between the new premise and the original premise.

4. Ensure that there is no direct association between the correct alternative and premise.
5. Avoid slang.

## B Amazon Mechanical Turk Form

### Task

#### 1. Read the following situation description:

*The customer filed a complaint with the store manager.*

#### 2. Tell us which of the following is the most likely CAUSE of this situation:

- (a) *The sales associate undercharged the customer.*
- (b) *The sales associate acted rude to the customer.*

Figure 3: Amazon Mechanical Turk task form

## C Example of an instance with low inter-annotator agreement

I received a package in the mail. What happened as a result? (effect)

- ✓ The package triggered my curiosity.
- ✗ I took the package to the post office.
- (a) Original COPA instance.

I received **someone's** package in the mail. What happened as a result? (effect)

- ✗ The package triggered my curiosity.
- ✓ I took the package to the post office.
- (b) Mirrored COPA instance.

Figure 4: An example of one of mirrored COPA instances with low inter-annotator agreement. Paying attention to the highlighted word is key to picking the correct alternative.