

DISCO: Distilling Counterfactuals with Large Language Models

Zeming Chen^{†*} Qiyue Gao^{‡*} Antoine Bosselut[†] Ashish Sabharwal[‡] Kyle Richardson[‡]

[†] Natural Language Processing Lab, EPFL, Lausanne, Switzerland

{zeming.chen, antoine.bosselut}@epfl.ch

[‡] Allen Institute for AI, Seattle, U.S.A.

{bertg, kyler, ashishs}@allenai.org

Abstract

Models trained with counterfactually augmented data learn representations of the causal structure of tasks, enabling robust generalization. However, high-quality counterfactual data is scarce for most tasks and not easily generated at scale. When crowdsourced, such data is typically limited in scale and diversity; when generated using supervised methods, it is computationally expensive to extend to new counterfactual dimensions. In this work, we introduce *DISCO* (**DIS**tilled **CO**unterfactual Data), a new method for automatically generating high-quality counterfactual data at scale. *DISCO* engineers prompts to generate phrasal perturbations with a large general language model. Then, a task-specific teacher model filters these generations to distill high-quality counterfactual data. While task-agnostic, we apply our pipeline to the task of natural language inference (NLI) and find that on challenging evaluations such as the NLI stress test, comparatively smaller student models trained with *DISCO*-generated counterfactuals are more robust (6% absolute) and generalize better across distributions (2%) compared to models trained without data augmentation. Furthermore, *DISCO*-augmented models are 10% more consistent between counterfactual pairs on three evaluation sets, demonstrating that *DISCO* augmentation enables models to more reliably learn causal representations. Our repository are available at: <https://github.com/eric1leca/disco>

1 Introduction

Despite the tremendous progress made in NLP on a wide range of reasoning tasks (Wang et al., 2018, 2019a; Xu et al., 2020), dataset biases continue to be a formidable challenge for robust model development (Gururangan et al., 2018; Poliak et al., 2018; Kaushik and Lipton, 2018; Tsuchiya, 2018; Liu et al., 2020b; Du et al., 2022). Counterfactual

* Work done while at the Allen Institute for AI. Equal contribution

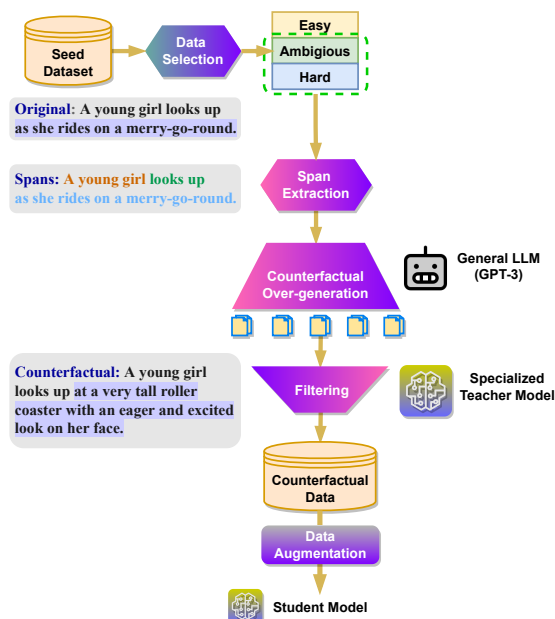


Figure 1: Overview of our counterfactual data distillation process (*DISCO*) using a large language model.

data augmentation (CAD) (Kaushik et al., 2019) is one general approach to improve model robustness by training on edited instances that systematically alter the critical or causally salient parts of dataset instances that contributes to the label assignment. To date, two main approaches have been pursued as part of these efforts: *human-centered approaches*, where edits are obtained through direct human annotation and crowdsourcing (Kaushik et al., 2019; Khashabi et al., 2020; Gardner et al., 2020); and *model-based approaches*, where new examples are collected through automatic text generation (Wu et al., 2021; Madaan et al., 2021; Ross et al., 2022; Wen et al., 2022, *inter alia*).

However, crowd-sourcing counterfactual data can be inefficient, costly, and difficult to scale. This often results in small counterfactual datasets, which can hinder the diversity and coverage of the collected edits (e.g., in Kaushik et al. (2019), the train-

ing scenario for NLI involves 8.3k total instances with augmentation). In contrast, supervised text generation methods are cheaper and easier to scale (e.g., Wu et al. (2022) use generation methods that expand NLP datasets to include around a million total examples). However, such methods can only generate fixed perturbation types. They rely on a fixed inventory of perturbation types each requiring new training sets. This is hard to scale up and can limit the space of perturbation types learned by the corresponding learned generation models. They can also be expensive to extend to new perturbation types, given the need to retrain models.

In this paper, we focus on the Natural Language Inference (NLI) task, which has recently been shown to benefit from collaboration between human annotation and LLMs in the WANLI data augmentation system of Liu et al. (2022). Our primary contribution is a *counterfactual knowledge distillation* procedure called *DISCO* (**DIS**tilled **CO**unterfactual Data), which works in the following way (see Figure 1): First, task instances to be edited are selected and decomposed into spans using off-the-shelf linguistic processing tools. Then prompt engineering and in-context learning are applied with a general LLM to overgenerate a diverse set of perturbations for these instances. We then employ a large *teacher* NLI model to conservatively filter the over-generations as a fully-automatic alternative to the human filtering used in WANLI. The distilled generations are finally used to train a much smaller and high-performance *student model*.

We show that *DISCO*, despite not relying on explicit human annotation, yields high-quality datasets. Manual annotation shows that, on average, 83% of our counterfactual data correctly flips the source labels, which is 1% higher than human performance. Additionally, compared to human CAD examples (Kaushik et al., 2019), we find *DISCO* generated data to have much-improved perturbation and information richness. Through data augmentation experiments, we also find that training on datasets built using *DISCO* obtains competitive and often improved performance across a wide range of robustness and out-of-domain (OOD) NLI tests, despite having a significantly smaller size than existing augmentation approaches (75k vs. 1 million from Wu et al. (2022)). This includes consistent improvements (6% average) over WANLI and SNLI baselines on 7 NLI robustness tests.

Building on the impressive results from Liu et al. (2022), this is significant as it shows the promising potential of data augmentation via LLMs, even without explicit human annotation. We find that models trained using our data exhibit 8% improved counterfactual accuracy and 6% increased sensitivity to context differences between counterfactual pairs than SNLI baselines. When augmenting on top of WANLI, our method shows an 18% performance gain on counterfactual accuracy.

Contributions In summary, we present *DISCO*, a fully-automatic counterfactual knowledge distillation approach based on LLMs. To our knowledge, *DISCO* is the first to use LLMs such as GPT3 for counterfactual data augmentation. We show that our approach helps produce more diverse counterfactuals over existing crowd-sourcing approaches while showing higher quality than human-written data. The distilled counterfactual data is more effective than existing augmentation approaches for improving NLI robustness, OOD generalization, and counterfactual consistency.

2 Related Work

Mitigating Spurious Correlations for NLU The augmentation methods described above are part of a large literature on model debiasing approaches, which also includes work on dataset filtering (Bras et al., 2020), model ensembling (Clark et al., 2019), feature removal, and other learning-based techniques (Belinkov et al., 2019; Mahabadi et al., 2020). Wu et al. (2022) also propose a new debiasing method called Z-Aug that learns to generate unbiased samples and filter out biased data using a z-statistic filter. In contrast to the debiasing and data generation techniques already discussed, our approach is unique in exploiting the power of LLMs such as GPT3 (Brown et al., 2020) to create more diverse augmented datasets as a way to mitigate biases and shortcuts.

Counterfactual Data Augmentation Augmenting models with counterfactual data is a popular recent approach for mitigating spurious correlation and improving model robustness. Kaushik et al. (2019) first recruits human workers to write counterfactual examples for augmentation. They find that counterfactually augmented data can help mitigate spurious patterns in the training data. As already discussed, however, creating counterfactual data using humans requires a high cost, is

time-consuming, and can result in simple perturbations. Later, Wu et al. (2021) and Ross et al. (2022) proposed frameworks that use text generation models to generate counterfactual data. These models require fine-tuning using pre-defined perturbation types. Both methods have constraints: (1) the generation is un-targeted, thus unlabeled, and (2) the perturbation types are limited. To acquire new perturbation types, the models have to be re-trained. Unlike the previous methods, our method uses LLMs to generate more diverse perturbation types cheaply and efficiently. Our method also improves over un-targeted generations by using a task-specific teacher model to verify the label.

Large Model Dataset Creation Leveraging the powerful generative ability of large language models to create datasets automatically has recently attracted considerable attention. This method reduces the cost of manually creating the dataset, can collect more diverse phenomena to expand the distribution, and can be adapted to a wide range of tasks in NLP. The most similar work to ours is WANLI (Liu et al., 2022), an NLI dataset fully generated by GPT-3 and annotated by human workers. The idea is to elicit ambiguous NLI examples from GPT-3 to improve its performance on challenge evaluation benchmarks, which relies on the *dataset cartography* techniques from Swayamdipta et al. (2020) that we also use in our study for selecting instances to edit. Our work also seeks to get diverse data from GPT-3 to improve model robustness. Differently, we only make local perturbations on the premise instead of generating a new example. We did not label our training data using human workers but leveraged an NLI model to filter out the counterfactual examples.

3 Counterfactual Distillation

The central idea of counterfactual data distillation is to prompt a large language model through in-context learning to generate perturbations that can flip the current label to a new one (ex. *Contradiction* \rightarrow *Entailment*). Once we select a subset of a dataset (discussed in Section 5.1), we first identify potential locations for performing counterfactual perturbations on the target instances. Then we prompt the GPT-3 (text-DaVinci-002) model to overgenerate perturbations (3.1). We use a teacher language model specializing in the NLI task to filter the generated perturbations based on the shift in model predictions from the orig-

inal to the new label (3.2). Formally, given an input premise-hypothesis pair $\langle P, H \rangle, l$ where $l \in \{Entailment, Contradiction, Neutral\}$ is the ground-truth label. We want to get a counterfactual input $\langle P', H \rangle, l'$ where we get P' by perturbing parts of the premise P and l' is the new label corresponding to the new input.

3.1 Prompting

We experiment with various prompting strategies on GPT-3, detailed and illustrated in Figure 2. To make local edits to a sentence following CAD (Kaushik et al., 2019)’s procedure, we use a neural syntactic parser (Akbik et al., 2019) to split sentences to perturb into spans. Using this neural chunker, we can get a set of spans $\mathcal{S} = \{s : s \in P\}$ decomposed from the premise P . These spans serve as the potential locations for making a perturbation.

Masked Prompting. To prompt GPT-3 for counterfactual perturbations, we use a masked NLI format to build the prompt. Let P and H be the premise and hypothesis pair we want to perturb, associated with the current label l and the set of spans \mathcal{S} . We select one span from \mathcal{S} and replace it in the premise with a mask token **[blank]**. Given a new label l' we want to flip to, we ask the model to fill in the blank mask token with creative perturbation s' to get a new premise P' that satisfies l' . Here the perturbation serves as an intervention in flipping the original label to the new label. Because during the generation time, one can not know which span will flip the label after perturbation, we overgenerate perturbations by iterating through all the spans from a premise. Each span yields a new prompt and makes a new request to GPT-3.

Insertion Mode. One of the key features of GPT-3 is its insertion mode, which allows users to insert a piece of text into the current context and have the model generate text based on the surrounding context. We can naturally convert the masked-NLI prompt into an insertion prompt format by providing the surrounding text of the mask token to the model. By forming a natural sentence, we try to align the prompt to the pre-training objective of GPT-3 (e.g., causal language modeling). We first map the label space $\{Entailment, Contradiction, Neutral\}$ to $\{true, false, possible\}$. Then we build the prompt: " \langle Prefix [insert] \rangle Suffix \rangle . It is $\langle l' \rangle$ that $\langle H \rangle$ ", where l' is the new label.

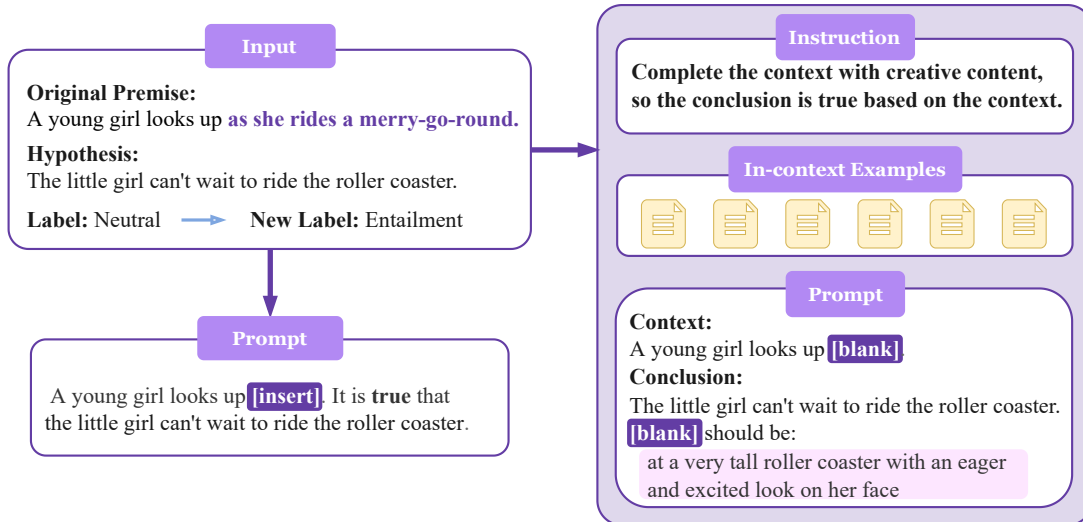


Figure 2: Overview of the perturbation generation process on GPT-3 using the masked NLI prompt (right) and the insertion NLI prompt (left-bottom). Here we are editing the input premise and hypothesis for the given NLI problem to change the label from *Neutral* to *Entailment*. Thus we have "conclusion is true" in both prompts.

The advantage of using the insertion mode is that the model considers both the prefix and suffix context of the masked span. This solves a common issue in the completion mode where the model tends to finish a sentence when generating the perturbation without noticing the suffix context. Additionally, the insertion mode does not require in-context learning examples, which yields more diverse generations at a much lower cost.

3.2 Teacher Model Filtering

Using a combination of the prompting strategies detailed in the last section, we then implement a filtering system to select the most promising counterfactual examples, pruning out potential mistakes made by GPT3. The filtering system first uses a heuristic-based automatic filter to remove any generations that yield obvious signs of low quality, ensuring that the remaining perturbations are more likely to flip the label in the desired direction. Our check for several criteria, including:

1. Does the perturbation contain parts from the instruction or prompt?
2. Does the perturbation copy parts from the in-context examples?
3. Does the perturbation repeat parts from the premise or hypothesis?

Using a count of the lexical overlap rate between sentences and a pre-defined set of common negation words, we also remove any perturbations with clear data artifacts, such as excessive lexical overlap between premise and hypothesis or using

negation words as a shortcut to flip the label. After the automatic filtering, we distill the remaining data using a model-based teacher, which identifies the perturbations that convert the original label to the target label. To verify if a perturbation potentially converts the original label in the direction of the new label, a natural way would be to check if the prediction probability of the new label shifts by a large margin, given the new input and the original input. Specifically, we calculate the distributional shift as follows:

$$\Delta_{l'} = p(l'|P', H) - p(l'|P, H), \quad (1)$$

which yields the change in prediction probability from the original input to the new input. We use a DeBERTa-v2 (He et al., 2020) model with SOTA performance on NLI as the teacher model. Additional details about the prompting parameters and teacher model can be found in Appendix A.

4 Evaluate Counterfactual Quality

Large general language models like GPT-3 enable the generation of counterfactual data at a large scale. The generation process is more efficient, cheap, and flexible than crowdsourcing. Here we evaluate the quality and diversity of *DISCO* data against counterfactually augmented data written by human workers (Human-CAD) (Kaushik et al., 2019) using automatic and human-based metrics.

Data	Flip Rate Score \uparrow						Avg.
	E2C	E2N	N2C	N2E	C2N	C2E	
Human-CAD	86.37	82.36	86.08	84.34	73.42	82.28	82.55
<i>DISCO</i> (ours)	78.53	82.70	76.20	85.53	75.76	92.43	83.14
	Soft Flip Rate Score \uparrow						Avg.
	E2C	E2N	N2C	N2E	C2N	C2E	
Human-CAD	94.32	83.33	88.61	86.75	82.28	94.94	88.24
<i>DISCO</i> (ours)	97.55	88.46	76.20	89.47	92.42	95.45	93.33
	Self-BLEU Diversity Score \downarrow						Avg.
	E2C	E2N	N2C	N2E	C2N	C2E	
Human-CAD	0.76	0.75	0.82	0.82	0.81	0.79	0.79
<i>DISCO</i> (ours)	0.23	0.26	0.26	0.18	0.25	0.21	0.23
	OTDD Dataset Distance \uparrow						Avg.
	E2C	E2N	N2C	N2E	C2N	C2E	
Human-CAD	217	95	179	139	238	217	180
<i>DISCO</i> (ours)	250	199	254	165	275	301	240

Table 1: Automatic and human evaluation results on a random subset (510 instances) of our counterfactual data (*DISCO*), compared with Human-CAD (Kaushik et al., 2019), counterfactual data written by human workers.

4.1 Automatic Evaluation

Diversity Measurement Following other work on CAD (Wu et al., 2021), we use Self-BLEU (Zhu et al., 2018) to measure the *diversity* of the generated counterfactual examples. In Table 1, we list the Self-BLEU score for each perturbation direction. Compared to human-written examples, GPT-3 generated examples have much lower Self-BLEU scores than human-written ones indicating that GPT-3 can generate far more diverse examples.

Dataset Distance The Self-BLEU score measures lexical and syntactic diversity only. To assess the diversity of information in the data, we calculate the dataset distance between the original examples and the new examples. Specifically, we measure dataset distance via OTDD (*optimal transport dataset distance*) (Alvarez-Melis and Fusi, 2020), a model-agnostic distance metric that can operate on datasets with disjoint label sets. OTDD can measure how well the knowledge from one dataset can transfer to another. We use OTDD to assess the distributional difference between the original and new examples. As Table 1 shows, our generated examples have a higher distance from the original examples than the human-written data, consistently in all directions. This trend demonstrates that our counterfactual data provide more diverse information than human-written data.

4.2 Human Evaluation

Label-Flip Score The label-flip score is an accuracy-based metric to check if the new exam-

ple after perturbation forms a counterfactual to the original example. We check the flip score in two aspects. The Label Flip Rate (LFR) calculates the percentage of new examples that flip the original label to the target label. The Soft Label Flip Rate (SLFR) calculates the percentage of new examples whose label differs from the original example’s label. SLFR measures how often LLMs generate valid counterfactuals independent of whether the new label is right. Given the rigidity of LFR and the fluidity of some NLI judgements (Pavlick and Kwiatkowski, 2019), this last metric is meaningful for checking if we still generate valid counterfactuals even when the exact label is not correct. The high SLFR suggests that many examples not accepted by the filter could be valid counterfactuals making them useful for other types of learning (e.g., leveraging signals from such data to train models to identify counterfactuals). For a dataset with K examples, we calculate FLR and SFLR as follows:

$$\text{LFR} = \frac{1}{K} \sum_{k=1}^K \mathbb{1}(\tilde{l}_k = l'_k)$$

$$\text{SLFR} = \frac{1}{K} \sum_{k=1}^K \mathbb{1}(\tilde{l}_k \neq l_k),$$

where \tilde{l} is the annotated label, l' is the target label, and l is the original label.

We use Amazon Mechanical Turk to conduct human evaluations, asking annotators to label a random subset of our data following the standard annotation process for the NLI task. We assigned three annotators for each example and did majority voting on the annotated labels. We list more details on the instructions, interface, and annotator requirements in Appendix B. We only give annotators the new sentence pairs to avoid bias from the original example. Table 1 shows the human evaluation results in each perturbation direction.

Compared to human-written examples, *DISCO* has lower LFRs only on generating contradictions, showing that GPT-3 generates better entailment and neutral examples rather than contradiction examples. We hypothesize that this is due to the ambiguous boundary between contradiction and neutral examples. Moreover, generating contradictions while maintaining diversity is difficult. When asked to generate contradictions, they tend to generate neutral examples by changing a sentence’s semantics (i.e., adding diversified words). In the case of Human-CAD, annotators tend to create con-

Dataset	Focus	Size
PI-CD (a)	Partial-input heuristics	3261
PI-SP (b)	Partial-input heuristics	371
IS-CS (c)	Inter-sentences Heuristics	656
LI-LI (d,e)	Logical Inference Ability	9927
LI-TS (f,g)	Logical Inference Ability	9832
ST (e)	Stress (distraction & noise) test	93447
HANS (h)	Syntactic Heuristic	30000
MNLI-hard-m	Out-of-distribution	4573
MNLI-hard-mm	Out-of-distribution	4530
QNLI	Out-of-distribution	5266
Human-CAD	Counterfactual consistency	1600
SNLI-hard $\square\rightarrow$	Counterfactual consistency	3042
WANLI $\square\rightarrow$	Counterfactual consistency	4000
(a) Gururangan et al. (2018)	(b) Liu et al. (2020a)	
(c) Nie et al. (2019)	(d) Glockner et al. (2018)	
(e) Naik et al. (2018)	(f) Minervini and Riedel (2018)	
(g) Wang et al. (2019b)	(h) McCoy et al. (2019)	

Table 2: Details about the evaluation datasets we used for the experiments.

tradictions using simple tricks like negation (Joshi and He, 2022). Although these tricks can produce absolute contradiction examples, they can introduce strong data artifacts, leading to a model that is not robust. Overall, the human evaluation scores show that our distilled counterfactual data exceeds human-written examples in correctly flipping the label, as shown by a higher average flip rate score.

5 Experiments

5.1 Counterfactual Data Augmentation

We next investigate how distilled GPT-3 counterfactual data can improve model robustness and generalizability through data augmentation. Given a set of original data $\mathcal{D} = \{\mathcal{X}, \mathcal{Y}\}$, we generate a perturbation z for each example in a subset of \mathcal{D} ($\mathcal{D}^s = \{\mathcal{X}^s, \mathcal{Y}^s\}$), and convert the original one to a counterfactual example: $\mathcal{D}^c = \{(x^c = z(x), y') | x \in \mathcal{X}^s, y \in \mathcal{Y}^s\}$. Next, we augment this subset by merging it with the counterfactual examples: $\mathcal{D}^a = \mathcal{D}^s \cup \mathcal{D}^c$. For additional data augmentation, we also select a base set \mathcal{D}^b (a random subset from \mathcal{D}), merge it with the augmentation set \mathcal{D}^a and remove any duplicated examples: $\mathcal{D}_{train} = \mathcal{D}^b \cup \mathcal{D}^a - \mathcal{D}^d$. We use models trained on base sets \mathcal{D}^b alone as baselines and evaluate whether augmenting the base sets using DISCO data would improve the baselines’ performances following Z-aug (Wu et al., 2022) and WANLI (Liu et al., 2022). We train a smaller student model, based on **RoBERTa-large** (355 million parameters) using the implementation from Wolf et al.

(2020), on \mathcal{D}_{train} and \mathcal{D}^a . Then, we evaluate the model on a set of test datasets for measuring robustness and OOD generalizability.

Source Datasets We select SNLI (Bowman et al., 2015) as the source dataset for generating *DISCO* data and for data augmentation. SNLI is a widely-used NLI dataset employed in numerous research studies. We apply data cartography (Swayamdipta et al., 2020) to select the ambiguous part of SNLI. The paper suggests that training on ambiguous data yields more robust models. Our intuition is that enhancing the ambiguous set with counterfactual examples would benefit the model’s learning. We also augment *DISCO* on WANLI (Liu et al., 2022) to analyze the benefits of counterfactual data augmentation on a dataset constructed via human-GPT-3 collaboration.

Evaluation Datasets We first evaluate how robust model performance is under adversarial and stress tests. We select the adversarial datasets from Liu et al. (2020b)’s benchmark for debiasing strategies and NLI stress test suite from Naik et al., 2018’s work. Next, we evaluate the model’s generalizability across different distributions. We select two datasets with a different distribution from the SNLI dataset: MNLI-hard (matched and mismatched) (Mahabadi et al., 2020), and QNLI (Wang et al., 2018), a dataset adapted from the Stanford Question Answering Dataset (Rajpurkar et al., 2016). Details about the evaluation datasets are included in Table 2.

Comparisons For naive comparison, we evaluate our models against baselines trained \mathcal{D}^b only without data augmentation. Then, we compare our models to prior augmentation methods, including Tailor (Ross et al., 2022), WANLI (Liu et al., 2022), Z-aug (Wu et al., 2022), and Human-CAD (Kaushik et al., 2019). For WANLI and Z-aug, we also augment them on the full SNLI training set because of their large dataset sizes. In addition, we fine-tune a model only on *DISCO* to compare with all the models above (see Appendix A for more details about training and hyper-parameters).

Results Table 3 shows that our counterfactual data augmentation significantly improves over the baseline performance on most robustness datasets when augmenting the *DISCO* dataset on a subset of SNLI. Augmenting or training with *DISCO* data achieves the highest accuracy on 7 evaluation

Method	Size	Model Robustness							OOD Generalization		
		PI-CD	PI-SP	IS-CS	LI-LI	LI-TS	ST	HANS	MNLI ¹	MNLI ²	QNLI
Large-size augmentation on full SNLI											
SNLI	549,367	82.2	69.0	68.4	93.6	72.5	72.4	73.1	78.5	78.2	64.5
+ WANLI	652,252	83.4	<u>82.7</u>	69.5	86.2	84.3	67.4	<u>87.4</u>	78.2	78.0	<u>78.6</u>
+ Z-aug	1,142,475	84.1	72.5	<u>72.6</u>	93.9	<u>87.1</u>	<u>75.4</u>	68.3	<u>80.0</u>	80.7	75.0
Augmentation on subset of SNLI											
SNLI-subset	100,000	82.0	71.7	65.1	85.5	83.9	69.5	65.8	78.0	79.1	73.4
+ Tailor	192,457	79.5	52.0	55.8	84.6	80.1	62.7	55.8	64.1	65.7	71.4
+ Human-CAD	108,330	82.8	<u>77.8</u>	69.2	90.7	87.1	71.3	65.5	79.0	79.0	72.8
+ <i>DISCO</i> (ours)	165,418	84.1	74.1	73.5	<u>92.1</u>	<u>88.4</u>	77.0	<u>70.1</u>	80.5	<u>80.2</u>	<u>77.7</u>
Augmentation on WANLI											
WANLI	102,885	65.6	81.3	65.9	65.6	82.7	56.5	89.4	76.1	76.3	81.1
+ <i>DISCO</i> (ours)	177,885	<u>82.8</u>	83.8	<u>72.0</u>	<u>86.8</u>	<u>85.1</u>	<u>68.6</u>	87.4	<u>80.0</u>	<u>78.7</u>	81.4
Trained on <i>DISCO</i> (ours) data only											
<i>DISCO</i> (ours)	75,000	83.5	77.4	73.3	89.4	88.9	76.3	70.7	79.2	79.5	79.1

Table 3: Results on Stress-tests, robust NLI test suites (Liu et al., 2020b), MNLI-hard, and QNLI. The bold numbers are the highest accuracy within a column, and the underlined numbers are the highest accuracy for each section. MNLI¹ refers to MNLI-hard-match, and MNLI² refers to MNLI-hard-mismatch.

sets. When augmenting on WANLI, the augmented model achieved better average performance (75.1) on robustness than the baseline WANLI model (65.9). We list the average performance gain for robustness and OOD generalization in Table 4. We can see that *DISCO*-augmented models improve model robustness over baselines by a large margin (6.5 SNLI and 9.5 WANLI). These results show the efficacy of our counterfactual data in helping models mitigate multiple types of NLI data bias altogether. On out-of-distribution (OOD) generalization, models trained on *DISCO* augmented data achieve a positive performance gain of 2.7% over the SNLI subset baseline and 2.1% over the WANLI baseline. This suggests that augmenting with *DISCO* helps the model generalize to datasets with distributional shifts. Compared to prior data augmentation methods, *DISCO* data can more significantly improve model performance, showing that our method yields high-quality and effective augmentation data.

In addition, *DISCO* is much smaller than other augmentation data like WANLI and Z-aug. Interestingly, training on *DISCO* data yields better performance than these models trained on large datasets (on 7 datasets).

5.2 Counterfactual Evaluation

In our second experiment, we investigate how *DISCO* data can enhance counterfactual reasoning ability of models on NLI problems. Counterfactual reasoning is the ability to predict how

	Test Metrics	Original	Augmented	Δ
SNLI-SUB	Robustness Avg.	71.0	77.5	6.5
	OOD Avg.	76.7	79.4	2.7
	Acc $_{\square \rightarrow}$ Avg.	47.1	55.2	8.1
	δ_s Avg.	58.6	64.9	6.3
WANLI	Robustness Avg.	65.9	75.1	9.2
	OOD Avg.	78.0	80.1	2.1
	Acc $_{\square \rightarrow}$ Avg.	34.6	52.7	18.1
	δ_s Avg.	44.9	57.6	12.7

Table 4: Performance gain of data augmentation using *DISCO* from baselines *without* augmentation (i.e., using the base sets in the first column).

an alternative context, contrary to the present context, might have resulted in different outcomes (Qin et al., 2019). In the setting of NLI, we alter the current context with text perturbations sufficient to change the current label to a different one while spuriously correlated features remain identical. A model that relies heavily on spurious features will likely fail to predict both the original and counterfactual examples correctly (Feder et al., 2022).

Evaluation Datasets We first create two counterfactual evaluation datasets using GPT-3 to generate the perturbations. We recruit human workers on Amazon Mechanical Turk to annotate labels for the two datasets. **SNLI-hard $_{\square \rightarrow}$** is constructed using a subset of the SNLI-hard (Gururangan et al., 2018) dataset. We pair each original example with the generated counterfactual example, where human annotators provide the gold label. In addi-

tion, we want to construct a dataset different from *DISCO*’s distribution. Thus, we select a subset from the WANLI test set and follow the same procedure as SNLI-hard $\square\rightarrow$ to get a counterfactual evaluation set WANLI $\square\rightarrow$. We assign three human workers to each problem to annotate the label. We list more details on the instructions, interface, and annotator requirements in Appendix B. We also include the **Human-CAD** dataset as the examples were written and labeled by human workers.

Metrics We measure models’ counterfactual reasoning ability along two dimensions. First, we measure *counterfactual sensitivity* δ_s : how confidently a model differentiates the original and counterfactual examples. In other words, how confidently does it assign a different label when there is a causal change in the input. Specifically, we define $\delta_s \in [0, 1]$ as:

$$\delta_s = \frac{(p(\hat{l}'|x') - p(\hat{l}'|x)) + (p(\hat{l}|x) - p(\hat{l}|x'))}{2},$$

where $x = (P, H)$ is the original input and x' is its perturbation. Intuitively, this metric quantifies the amount of shift in model predictions between the two related examples. Unchanged model prediction results in a sensitivity of 0. When model prediction changes with extremely high confidence (i.e., assigning 100% on its predicted labels), δ_s is normalized to 1. In binary classification, when the predicted label changes, the metric simplifies to:

$$\delta_s = p(\hat{l}'|x') + p(\hat{l}|x) - 1.$$

δ_s here measures the model’s confidence in prediction when the context changes, shown by the probability it assigns to the predicted labels. In general, the higher the δ_s , the more sensitive the model is to context changes in the input.

Next, we measure the counterfactual accuracy $\text{Acc}_{\square\rightarrow}$. Under this metric, a prediction is correct only when the model correctly predicts the original and counterfactual examples. We use counterfactual accuracy to measure the consistency of model performance on original and counterfactual examples. $\text{Acc}_{\square\rightarrow}$ is defined as:

$$\frac{1}{K} \sum_{k=1}^K \mathbb{1} \left((\hat{l}_k | P_k, H_k) = l_k^* \wedge (\hat{l}'_k | P'_k, H_k) = l'^*_k \right),$$

where K is the number of examples evaluated, \hat{l}, \hat{l}' are model predictions for the original and counterfactual examples, and l^*, l'^* are the gold labels,

Method	Human-CAD		SNLI-hard $\square\rightarrow$		WANLI $\square\rightarrow$	
	δ_s	$\text{Acc}_{\square\rightarrow}$	δ_s	$\text{Acc}_{\square\rightarrow}$	δ_s	$\text{Acc}_{\square\rightarrow}$
SNLI-subset	62.8	59.1	66.1	51.1	51.3	39.3
+ Tailor	58.8	55.6	60.6	55.6	33.9	23.7
+ Human-CAD	70.9	63.6	73.6	54.1	34.6	42.8
+ <i>DISCO</i> (ours)	69.4	64.1	74.3	60.3	55.9	47.7
WANLI	41.4	30.5	47.4	27.0	44.5	42.1
+ <i>DISCO</i> (ours)	65.6	64.9	68.5	59.2	46.1	42.8
<i>DISCO</i> (ours)	65.7	66.5	71.2	63.1	41.9	48.3

Table 5: Performance on the counterfactual sensitivity tests including three datasets: Human-CAD, SNLI-hard $\square\rightarrow$, AND WANLI $\square\rightarrow$. The models used for evaluation are from the first experiment directly.

respectively. This is similar in spirit to evaluations based on *contrast sets* from Gardner et al. (2020), *perturbation clusters* from Khashabi et al. (2020), and the *grouped probe metric* of Trivedi et al. (2020).

Results Table 5 shows models’ performance on the three counterfactual evaluation sets. Models augmented or trained with *DISCO* consistently outperform the baseline models by a large margin. Training with only *DISCO* achieves the highest counter accuracy while augmenting *DISCO* on the SNLI subset achieves the highest counterfactual sensitivity. This shows that our data helps increase the model’s ability to differentiate the two examples and improve its reasoning performance on counterfactual data. Compared to other data augmentation methods, *DISCO* yields a performance gain on both metrics showing its benefit on counterfactual reasoning.

DISCO increases the WANLI baseline’s sensitivity and accuracy by more than 20% and 30% respectively on both Human-CAD and SNLI-hard $\square\rightarrow$. However, the increase on WANLI $\square\rightarrow$ is marginal, which is likely because *DISCO* and the WANLI train set have very different distributions (OTDD distance 579). Although WANLI $\square\rightarrow$ is close to the WANLI train set (OTDD distance 270), training on it yields lower accuracy than *DISCO*, indicating that human-GPT-3 collaborated data construction does not necessarily grant models the ability to reason on counterfactual data. Thus, we can confirm that the distillation step on top of GPT-3 generation is essential for improving the model’s counterfactual reasoning ability.

6 Conclusion

In this paper, we introduced the *DISCO* framework for distilling high-quality counterfactual data from large language models (LLMs) using a task-specific teacher model for NLI. Through automatic and human evaluations, we show that counterfactuals generated by LLMs have higher quality and accuracy than human-written examples while having more diverse perturbations. Our evaluation results suggest that training or augmenting with distilled counterfactual data can help mitigate various types of distinct spurious patterns. Counterfactual examples produced by *DISCO* significantly benefit model performance with improved robustness and out-of-distribution (OOD) generalizability. Despite a smaller data size, *DISCO* data help models achieve better performance on the evaluation sets than baselines with extensive data. Furthermore, training on *DISCO* examples improves model performance on counterfactual accuracy and helps the model be more sensitive to the context changes between counterfactual and original examples.

For future work, our method suggests several directions. While our efforts are limited to NLI, generating counterfactual data using LLMs is more general and, we believe, can be fruitfully applied to a wider range of tasks. In specific, only a task-specific filter model and modification to LLM prompts are needed to extend our generation pipeline to other tasks or even other languages. Also, while our approach takes inspiration from knowledge distillation (Hinton et al., 2015) approaches and relies on a *teacher* filtering model, alternative strategies could be used to improve the quality. As a related direction, techniques for semi-supervised learning over unfiltered LLM output should also be investigated to help utilize the wide range of data produced by LLMs.

7 Limitations

While we have argued that our approach to collecting counterfactual data via *DISCO* is agnostic to the particular task and language, we emphasize that the experiments we report are limited to English and the task of NLI. Given that English is a high-resource language, there could be additional challenges (e.g., finding the tools needed for making span selection) in re-creating our pipeline for other languages. We also emphasize that our data generation experiments were carried out using only a single LLM, namely the publicly available GPT3

model first reported in Brown et al. (2020).

As with the related studies we cite (e.g., Liu et al. (2022)), given the high costs associated with large-scale prompting, we are unable to ablate all parts of our data generation pipeline (e.g., the effect of systematically alternating prompting styles at scale, alternative span extraction techniques). Similar to virtually all experiments involving LLM prompting, such differences could affect the results and quality of the resulting augmentation datasets. Similarly, given the high costs of human annotation, we have limited our human evaluation to around 500 random instances (each involving 3 annotators), which follows other related studies.

Acknowledgements

We thank the anonymous reviewers for their constructive and thoughtful comments. We also thank the members of the Aristo team at AI2 for providing helpful feedback on earlier versions of this work. Thanks finally to the beaker.org team at AI2 for their assistance and help with experiments and computing infrastructure.

References

- Alan Akbik, Tanja Bergmann, Duncan Blythe, Kashif Rasul, Stefan Schweter, and Roland Vollgraf. 2019. *FLAIR: An easy-to-use framework for state-of-the-art NLP*. In *Proceedings of NAACL*.
- David Alvarez-Melis and Nicolò Fusi. 2020. *Geometric dataset distances via optimal transport*. In *Proceedings of NeurIPS*.
- Yonatan Belinkov, Adam Poliak, Stuart M Shieber, Benjamin Van Durme, and Alexander M Rush. 2019. *Don't take the premise for granted: Mitigating artifacts in natural language inference*. *Proceedings of ACL*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. *A large annotated corpus for learning natural language inference*. In *Proceedings of EMNLP*.
- Ronan Le Bras, Swabha Swayamdipta, Chandra Bhagavatula, Rowan Zellers, Matthew E. Peters, Ashish Sabharwal, and Yejin Choi. 2020. *Adversarial filters of dataset biases*. *Proceedings of ICML*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. *Language models are few-shot learners*. *Proceedings of NeurIPS*.

- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. [Don't take the easy way out: Ensemble based methods for avoiding known dataset biases](#). *Proceedings of EMNLP*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of NAACL*.
- Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. 2022. [Shortcut learning of large language models in natural language understanding: A survey](#).
- Amir Feder, Katherine A. Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E. Roberts, Brandon M. Stewart, Victor Veitch, and Diyi Yang. 2022. [Causal inference in natural language processing: Estimation, prediction, interpretation and beyond](#). *Transactions of the Association for Computational Linguistics*, 10:1138–1158.
- Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, et al. 2020. [Evaluating models' local decision boundaries via contrast sets](#). *Findings of EMNLP*.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking nli systems with sentences that require simple lexical inferences](#). *Proceedings of ACL*.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. 2018. [Annotation artifacts in natural language inference data](#). *Proceedings of NAACL*.
- Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2020. [Deberta: Decoding-enhanced bert with disentangled attention](#).
- Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. 2015. [Distilling the knowledge in a neural network](#). *arXiv preprint arXiv:1503.02531*, 2(7).
- Nitish Joshi and He He. 2022. [An investigation of the \(in\)effectiveness of counterfactually augmented data](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3668–3681, Dublin, Ireland. Association for Computational Linguistics.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2019. [Learning the difference that makes a difference with counterfactually-augmented data](#). *Proceedings of ICLR*.
- Divyansh Kaushik and Zachary C. Lipton. 2018. [How much reading does reading comprehension require? a critical investigation of popular benchmarks](#). In *Proceedings of EMNLP*.
- Daniel Khashabi, Tushar Khot, and Ashish Sabharwal. 2020. [More bang for your buck: Natural perturbation for robust question answering](#). *Proceedings of EMNLP*.
- Diederik P Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). *Proceedings of ICLR*.
- Alisa Liu, Swabha Swayamdipta, Noah A Smith, and Yejin Choi. 2022. [WANLI: Worker and AI collaboration for natural language inference dataset creation](#). *Findings of EMNLP*.
- Tianyu Liu, Zheng Xin, Baobao Chang, and Zhifang Sui. 2020a. [HypoNLI: Exploring the artificial patterns of hypothesis-only bias in natural language inference](#). In *Proceedings of LREC*.
- Tianyu Liu, Zheng Xin, Xiaoan Ding, Baobao Chang, and Zhifang Sui. 2020b. [An empirical study on model-agnostic debiasing strategies for robust natural language inference](#). In *Proceedings of CoNLL*.
- Nishtha Madaan, Inkit Padhi, Naveen Panwar, and Dip-tikalyan Saha. 2021. [Generate your counterfactuals: Towards controlled counterfactual generation for text](#). In *Proceedings of AAAI*.
- Rabeeh Karimi Mahabadi, Yonatan Belinkov, and James Henderson. 2020. [End-to-end bias mitigation by modelling biases in corpora](#). *Proceedings of ACL*.
- R Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference](#). *Proceedings of ACL*.
- Pasquale Minervini and Sebastian Riedel. 2018. [Adversarially regularising neural nli models to integrate logical background knowledge](#). *Proceedings of CoNLL*.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. [Stress test evaluation for natural language inference](#). *Proceedings of COLING*.
- Yixin Nie, Yicheng Wang, and Mohit Bansal. 2019. [Analyzing compositionality-sensitivity of nli models](#). *Proceedings of the AAAI*.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of ACL*.
- Ellie Pavlick and Tom Kwiatkowski. 2019. [Inherent disagreements in human textual inferences](#). *Transactions of the Association for Computational Linguistics*, 7:677–694.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. [Hypothesis only baselines in natural language inference](#). In *Proceedings of *SEM*.
- Lianhui Qin, Antoine Bosselut, Ari Holtzman, Chandra Bhagavatula, Elizabeth Clark, and Yejin Choi. 2019. [Counterfactual story reasoning and generation](#). In *Proceedings of EMNLP*.

- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). *Proceedings of EMNLP*.
- Alexis Ross, Tongshuang Wu, Hao Peng, Matthew E Peters, and Matt Gardner. 2022. [Tailor: Generating and perturbing text with semantic controls](#). *Proceedings of ACL*.
- Swabha Swayamdipta, Roy Schwartz, Nicholas Lourie, Yizhong Wang, Hannaneh Hajishirzi, Noah A Smith, and Yejin Choi. 2020. [Dataset cartography: Mapping and diagnosing datasets with training dynamics](#). *Proceedings of EMNLP*.
- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. [Fever: a large-scale dataset for fact extraction and verification](#). *Proceedings of NAACL*.
- Harsh Trivedi, Niranjan Balasubramanian, Tushar Khot, and Ashish Sabharwal. 2020. [Is multihop QA in DiRe condition? Measuring and reducing disconnected reasoning](#). In *Proceedings of EMNLP*.
- Masatoshi Tsuchiya. 2018. [Performance impact caused by hidden bias of training data for recognizing textual entailment](#). *Proceedings of LREC*.
- Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2019a. [Superglue: A stickier benchmark for general-purpose language understanding systems](#). *Proceedings of NeurIPS*, 32.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2018. [Glue: A multi-task benchmark and analysis platform for natural language understanding](#). *Proceedings of ICLR*.
- Haohan Wang, Da Sun, and Eric P. Xing. 2019b. [What if we simply swap the two text fragments? a straightforward yet effective way to test the robustness of methods to confounding signals in nature language inference tasks](#).
- Jiaxin Wen, Yeshuang Zhu, Jinchao Zhang, Jie Zhou, and Minlie Huang. 2022. [Autocad: Automatically generating counterfactuals for mitigating shortcut learning](#). *Proceedings of EMNLP Findings*.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2018. [A broad-coverage challenge corpus for sentence understanding through inference](#). *Proceedings of NAACL*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of EMNLP*.
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel S Weld. 2021. [Polyjuice: Generating counterfactuals for explaining, evaluating, and improving models](#). *Proceedings of ACL*.
- Yuxiang Wu, Matt Gardner, Pontus Stenetorp, and Pradeep Dasigi. 2022. [Generating data to mitigate spurious correlations in natural language inference datasets](#). *Proceedings of ACL*.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, Yin Tian, Qianqian Dong, Weitang Liu, Bo Shi, Yiming Cui, Junyi Li, Jun Zeng, Rongzhao Wang, Weijian Xie, Yanting Li, Yina Patterson, Zuoyu Tian, Yiwen Zhang, He Zhou, Shaowei Hua Liu, Zhe Zhao, Qipeng Zhao, Cong Yue, Xinrui Zhang, Zhengliang Yang, Kyle Richardson, and Zhenzhong Lan. 2020. [CLUE: A Chinese language understanding evaluation benchmark](#). In *Proceedings of COLING*.
- Yaoming Zhu, Sidi Lu, Lei Zheng, Jiaxian Guo, Weinan Zhang, Jun Wang, and Yong Yu. 2018. [Tegygen: A benchmarking platform for text generation models](#).

A Hyper-parameters and Implementation

GPT3 and Teacher Model For perturbation overgeneration, we use GPT-3 with the text-DaVinci-002 version. We set the *temperature* to 0.8 to encourage creative generations. For the penalties, we set the *frequency penalty* and *presence penalty* to 0.8 to lower the likelihood of sampling repeated words. To mitigate error propagation from the filtering step, we use a publicly available DeBERTa-v2 (He et al., 2020) model checkpoint (containing 1.3 billion parameters) trained on a mixture of NLI datasets, including SNLI (Bowman et al., 2015), MultiNLI (Williams et al., 2018), FEVER (Thorne et al., 2018), ANLI (Nie et al., 2020), that achieves SOTA performance on these datasets.

Student Models and Training Protocol For all experiments, we tuned Robert-large (containing 345 million parameters) via a random search over key hyper-parameters in the style of Devlin et al. (2019). We used ADAM (Kingma and Ba, 2015) as our optimizer. The key hyper-parameters include *learning rate* (including $2e - 5$, $3e - 5$, $5e - 5$), *batch size* (between 32, 64), *warmup ratio* (in the range of 0.08, 0.1) and *number of epochs* (3 to 5); weight decay was kept constant at 0.1 following Liu et al. (2022), and early stopping was used with a patience of 2 epochs. We generally found the following configuration to yield good performance: LR= $3e - 5$, epochs=3, batch_size=64, warmup_ratio=0.1. Standardly, model selection was performed by choosing the model with the highest validation accuracy. In our main result tables (i.e., Tables 3-4) we report the best of 5 models based on random restarts with different random seeds in all rows excluding the first 3. In the first 3 rows, given the large size of the training sets and the generally high cost of fine-tuning, we report the best single run (and generally found these models to yield low variance across hyper-parameters).

When comparing against other data augmentation approaches, e.g., Z-aug (Wu et al., 2022), we used the exact code base compared with models trained on *DISCO* to remove any differences in implementation (our implementation is based on the transformers library (Wolf et al., 2020)). All experiments were performed on an NVIDIA RTX A6000 GPU.

B Human Annotation Details

We recruit human annotators to evaluate our generated counterfactual data and to annotate two evaluation sets for counterfactual consistency: SNLI-hard $\square\rightarrow$ and WANLI $\square\rightarrow$. Here we discuss the details of our annotation studies. Screenshots of the instructions, guidelines, and annotation interface are shown in Fig 3 and Fig 4.

Annotators We recruit human workers on the Amazon Mechanical Turk¹ platform. We required Mechanical Turk Masters to perform our tasks. Annotators must have a HIT approval rate of 98%, a total of 1000 approved HITs, and be in the United States. Throughout the data collection process, we randomly select a subset of the annotations to check and correct any potentially controversial annotations. For each problem, we assign three annotators and use a majority vote to determine the final annotation. Workers were paid \$0.3 for each AMT hit (consisting of 10 examples to annotate).

¹<https://www.mturk.com/>

Determine if the relationship between sentence A and sentence B, whether sentence A **entails** sentence B (Given Sentence A, you are confident that Sentence B is true or would happen) or sentence A **contradicts** with sentence B (Given Sentence A, you are confident that Sentence B cannot happen or is false) or **neither** (can not be determined).

<p>Sentence A: The ponies look like they love each other. Sentence B: The ponies are being intimate. Answer: Entails. Explanation: <u>Loving each other means they are being intimate.</u></p>	<p>Sentence A: A man with glasses and a red vest operates a machine. Sentence B: A man is at cafe Answer: Cannot be determined. Explanation: <u>The man might be an employee in the cafe, operating a coffee machine.</u></p>	<p>Sentence A: City workers are doing repairs. Sentence B: The workers are not working today. Answer: Contradicts. Explanation: <u>Doing repairs means the workers are working today.</u></p>
<p>Sentence A: A kid is skateboarding on a handrail. Sentence B: The kid does a skateboarding trick. Answer: Entails. Explanation: <u>Skateboarding on a handrail is a cool way of skateboarding.</u></p>	<p>Sentence A: A couple walking down the street. Sentence B: People getting exercise. Answer: Cannot be determined. Explanation: <u>It's unsure whether walking serves as this couple's exercise.</u></p>	<p>Sentence A: A woman in pink is covering her face with a scarf. Sentence B: A woman puts a scarf on the girl. Answer: Contradicts. Explanation: <u>The woman is putting the scarf on her face instead of on the girl.</u></p>

Figure 3: The annotated examples with explanations used on Amazon Mechanical Turk.

Instructions:
 Determine if the relationship between sentence A and sentence B, whether sentence A **entails** sentence B (Given Sentence A, you are confident that Sentence B is true or would happen) or sentence A **contradicts** with sentence B (Given Sentence A, you are confident that Sentence B cannot happen or is false) or **neither** (can not be determined).

Sentence A: Blond woman in a white shirt is cleaning a floor with a mop.
Sentence B: A woman is cleaning the floor.

Entails
 Contradicts
 Cannot be determined

Sentence A: Two people are sitting at a table in a restaurant on a city street.
Sentence B: Two people are sitting in a restaurant in the city.

Entails
 Contradicts
 Cannot be determined

Figure 4: Instructions provided to human annotators on Amazon Mechanical Turk and the annotation interface.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 7
- A2. Did you discuss any potential risks of your work?
Section 8
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract, Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3, 4, 5.1, 5.2

- B1. Did you cite the creators of artifacts you used?
Section 3, 4, 5.1, 5.2
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Not applicable. Left blank.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Table 2, 3

C Did you run computational experiments?

Section 5.1, 5.2

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Appendix A

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?
Section 5.1, 5.2; Appendix A
 - C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?
Section 5.1, 5.2
 - C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?
Section 3.1
- D** **Did you use human annotators (e.g., crowdworkers) or research with human participants?**
Section 4.2, 5.2
- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?
Appendix B
 - D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?
Appendix B
 - D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?
Not applicable. Left blank.
 - D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?
Not applicable. Left blank.
 - D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?
Not applicable. Left blank.