Plug and Play Counterfactual Text Generation for Model Robustness

***Abstract*—Generating counterfactual test-cases is an important backbone for testing NLP models and making them as robust and reliable as traditional software. In generating the test-cases, a desired property is the ability to control the test-case generation in a flexible manner to test for a large variety of failure cases and to explain and repair them in a targeted manner. In this direction, significant progress has been made in the prior works by manually writing rules for generating controlled counterfactuals. However, this approach requires heavy manual supervision and lacks the flexibility to easily introduce new controls. Motivated by the impressive flexibility of the plug-and-play approach of PPLM, we propose bringing the framework of plug-and-play to counterfactual test case generation task. We introduce CASPer, a plug-and-play counterfactual generation framework to generate test cases that satisfy goal attributes on demand. Our plug-and-play model can steer the test case generation process given any attribute model without requiring attribute-specific training of the model. In experiments, we show that CASPer effectively generates counterfactual text that follow the steering provided by an attribute model while also being fluent, diverse and preserving the original content. We also show that the generated counterfactuals from CASPer can be used for augmenting the training data and thereby fixing and making the test model more robust.**

I. Introduction

Machine learning and deep learning-based decision making has become part of today's software. This creates the need to ensure that machine learning and deep learning-based systems are as trusted as traditional software with increased deployment and wider-use. Traditional software is made dependable by following rigorous practice like static analysis, testing, debugging, verifying, and repairing throughout the development and maintenance life-cycle. Similarly, for testing and repairing NLP systems, we need inputs where models can fail and thereby bringing out issues early on [52], [53]. For this, counterfactual text data [12], [62] can be used. By treating counterfactual text as test cases, we are asking: *Would the model fail if the input text was modified to have different characteristics*? Furthermore, with such counterfactual text, NLP systems can be repaired by augmenting the training samples with these counterfactual test cases and its labels [10]. Hence, enabling model repair by generating counterfactual text is a crucial step in deploying these NLP systems more widely.

An important aspect of model testing and repair is to ensure that we can *control* these counterfactual test cases. The ability to control will allow us to test for specific types of failures that are important for the deployed model. Controlled counterfactuals can also allow us to fix the failures by creating new training samples in a focused manner for augmenting the existing training dataset. Thus we require a model that can generate counterfactuals that can be controlled by providing some goal attributes.

In addition to controlling the test-cases, we would also like to have *flexibility* about which goal attributes to apply and the flexibility to chain together multiple goal attributes in order to test how the deployed model behaves for a wide variety of textual characteristics. Bringing such flexibility requires a model that allows us to plug-and-play new attribute goals as and when required.

In this work, we propose a framework for counterfactual test-case generation also called *Counterfactual Sentence Generation with Plug-and-Play Perturbation* or CASPer that provides both control and flexibility during test case generation. To achieve these, we build on the framework of Plug-and-Play Language Models or PPLMs [9]. PPLMs have shown an impressive ability to flexibly steer pre-trained language models to generate attribute-conditioned text samples. However, PPLMs cannot be directly used for perturbing an input text and generating counterfactual samples. In this work, instead of steering language models, we steer a text-to-text model that is pre-trained to reconstruct its text inputs and we steer this model in a plug-and-play manner. Thus, like PPLM, our model is plug-and-play and is capable of generating counterfactuals for any arbitrary goal attributes provided at sampling time. In CASPer, we take BART [151] as our text-to-text reconstruction model. To generate each counterfactual, we perturb the hidden layer of the BART model similarly to PPLM to sample the counterfactual text. In experiments, we apply our framework to generate counterfactuals by providing named-entity and sentiment-based goal attributes. We show that our simple plug-and-play framework can generate counterfactuals that are fluent and content-preserving while also attaining the goal attributes and being effective as training samples for data-augmentation to improve the performance of the deployed model. We show counterfactual samples generated by CASPer and from existing models in Table I.

The main contribution of the paper can be seen as three folds: 1) We propose, CASPer, the first plug-and-play counterfactual generation model that achieves both control and flexibility. 2) Empirically, we show that our approach can generate fluent counterfactuals that preserve the content and also attain the goal attribute. 3) We also show the effectiveness of our counterfactuals as new training data in making the test models robust.

Inputs	**Token-based** [15]	**Adversarial** [42]	**Polyjuice** [26]	**Ours**
Input Text: Me and a group of friends rent horrible videos to laugh at them, trust me it has lead to some horribly spent money but also some great laughs. *Initial State* : No location named-entity is present in the text. *Steering Goal*: To make the sentence contain at least one **location** named-entity.	Me and a group of friends rent **youtube** videos to laugh at them, trust me it has lead to some horribly spent money but also some great laughs. Me and a group of friends rent **music** videos to laugh at them, trust me it has lead to some horribly spent money but also some great laughs.	Me and a group of friends rent horrible videos to laugh at them , trust me it has lead to some horribly spent money but also some flavorful laughs .	Me and a group of my friends rent horrible videos to laugh at them , trust me it has lead to some horribly spent money but also some **flavorful** laughs. and a group of friends rent videos to laugh at them, trust me it has lead to some horribly spent money but also some **chuckles**.	Me have a group of lads in **Brisbane** and rent horrible videos to get great laughs at. Some extremely expensive videos but some very great.. Me and a group of **Fairfax Bay** friends like to rent videos for laughs. Have spent some money but had great laughs at their terrible videos.

TABLE I

OVERVIEW OF GENERATIONS FROM THE EXISTING MODELS. WE PROVIDE A TEXT AS INPUT TO THE MODEL WITH A STEERING GOAL OF INTRODUCING A *location named-entity* INTO THE GIVEN TEXT. WE SHOW THE OUTPUTS FROM A TOKEN-BASED SUBSTITUTION MODEL [15], FROM ADVERSARIAL GENERATION [42], FROM POLYJUICE [26] AND FROM OUR PROPOSED MODEL. WE NOTE THAT TOKEN-BASED SUBSTITUTION METHOD, RELYING ON TEMPLATE MATCHING, FAIL TO MATCH A TEMPLATE AND ARE THUS NOT ABLE TO ACHIEVE THE STEERING GOAL. ADVERSARIAL, DUE TO ITS GRADIENT-DESCENT-BASED TOKEN-SUBSTITUTION, FAILS TO GENERATE PLAUSIBLE TEXT. POLYJUICE, DUE TO ITS TEMPLATE MATCHING, CHANGES VERY INSIGNIFICANT PART OF THE TEXT. OUR MODEL, TAKING ADVANTAGE OF BART AUTO-ENCODER, EFFECTIVELY ACHIEVES THE STEERING GOAL.

II. PRELIMINARIES

A. Counterfactual Text Generation for Model Testing and Repair

Taking a text $\mathbf{x}$ from the input distribution and modifying it $\mathbf{x} \rightarrow \mathbf{y}$ is known as the task of counterfactual text generation. However, our goal of counterfactual text generation is to help improve NLP models by using counterfactual text as test-cases in various stages of the model deployment. For such models, we would like to *i)* test for failures, *ii)* explain when those failures occur and *iii)* fix the failures by augmenting the training data with new training samples.

However, from the perspective of model repair, it is not enough to simply obtain uncontrolled and random perturbations $\mathbf{x} \rightarrow \mathbf{y}$ to generate these test-cases. We would like these test-cases to be controlled $\mathbf{x} \xrightarrow{\texttt{control}} \mathbf{y}$ through a given `control` input. By controlling the generated test cases, we would be able to *i)* test for specific types of failures that are important for the deployed model, *ii)* explain which controls lead to high failure and *iii)* create targeted data sets for augmenting the training set and fixing the models. Hence, our work lies at the intersection of counterfactual text generation and controlled text generation. In the following subsection, we provide an overview of controlled text generation.

B. Controlled Text Generation

The goal of controlled text generation is to generate samples $\mathbf{x}$ from a controlled distribution $p(\mathbf{y}|\mathbf{a})$ which is conditioned on a specific attribute or control $\mathbf{a}$. For example, the language model $p(\mathbf{y}|\mathbf{a})$ may be used to generate product reviews conditioned on a specific product category by setting $\mathbf{a} = \texttt{kitchen}$.

1) Plug-and-Play Language Models: Plug-and-Play Language Models (or simply PPLMs) provide an attractive solution to model the class-conditional distribution $p_\theta(\mathbf{y}|\mathbf{a})$. Plug-and-play models take a pre-trained unconditional generative model $p_\theta(\mathbf{y})$ and use the reward signal from the attribute model $p(\mathbf{a}|\mathbf{y})$ to quickly (in $\sim$10 gradient steps) modify the unconditional generative model $p_\theta(\mathbf{y})$ to generate samples from the desired distribution $p_\theta(\mathbf{y}|\mathbf{a})$.

To achieve this, PPLMs [9] take GPT-2 to be the unconditional generative model $p_\theta(\mathbf{y})$. In GPT-2, the text generation is done iteratively word-by-word. In each iteration t, one word is predicted and is fed back to the Transformer to predict the next word. This generation process can be described as follows:

$$
\begin{aligned}
H_t &= \text{Transformer}(\mathbf{y}_{<t}), \\
\mathbf{o}_t &= \text{PredictionHead}(H_t), \\
\mathbf{y}_t &\sim \text{Categorical}(\mathbf{o}_t).
\end{aligned}
$$

where t is the word position in the text, H_t is the last hidden layer before the prediction head and $\mathbf{o}_t$ are the log-probabilities of the words in the vocabulary used for sampling the next word $\mathbf{y}_t$. We shall refer to this model as the unmodified language model and denote the distribution that it models for the next word prediction as $p(\mathbf{y}_t|\mathbf{y}_{<t})$.

To generate a text from $p_\theta(\mathbf{y}|\mathbf{a})$ at test time, PPLMs learn a perturbation for the hidden state H_t of the unconditional

model $p_\theta(\mathbf{y})$. This is achieved as follows:

$$\begin{aligned} H_t &= \text{Transformer}(\mathbf{y}_{<t}), \\ \mathbf{o}_t &= \text{PredictionHead}(H_t + \Delta H_t), \\ \mathbf{y}_t &\sim \text{Categorical}(\mathbf{o}_t). \end{aligned}$$

where ΔH_t is the learned perturbation. We shall refer to this model as the modified language model and denote the distribution that it models for the next word prediction as $\bar{p}(\mathbf{y}_t|\mathbf{y}_{<t})$. The learning of the perturbation parameters $\{\Delta H_1, \ldots, \Delta H_T\}$ is driven by the following objective:

$$\begin{aligned} \mathcal{L}_{\text{PPLM}} = &-\log p(\mathbf{a}|\mathbf{y}) \\ &- \sum_{t=1}^{T} D_{\text{KL}}(p(\mathbf{y}_t|\mathbf{y}_{<t})||\bar{p}(\mathbf{y}_t|\mathbf{y}_{<t})), \end{aligned}$$

where the first term provides the learning signal to steer the generation towards the desired class or attribute by trying to maximize the log-probability of the desired attribute. The second term tries to keep the generations close to the unmodified language model to ensure that the text remains fluent and plausible. We note that this learning process is done separately each time we need to generate a new sample. However, the learning of the perturbation parameters $\{\Delta H_1, \ldots, \Delta H_T\}$ can be done very quickly and it only takes about 10 gradient steps. This property makes PPLMs flexible during generation.

PPLMs provide some useful properties that are lacking in other conditional generative models that are not plug-and-play. Plug-and-play models are *flexible* during sampling – meaning that new class-conditioning can be easily introduced at test time by simply replacing the attribute model $p(\mathbf{a}|\mathbf{y})$ with a new attribute model for the new class without requiring costly retraining with respect to the new attribute. Furthermore, plug-and-play models can support conditioning on logical clauses by simply composing multiple attribute models together. For instance, to generate product review text conditioned on a logical clause `kitchen + electronics + not electrical`, the attribute model $p(\mathbf{a}|\mathbf{y})$ can be written as the product of individual attribute models $p(\texttt{kitchen}|\mathbf{y}) \cdot p(\texttt{electronics}|\mathbf{y}) \cdot [1 - p(\texttt{electrical}|\mathbf{y})]$ and easily plugged into the PPLM framework [9]. Our model architecture is shown in Figure 1.

III. Method: Plug-and-Play Counterfactual Generation Framework

We now describe our proposed method to generate controlled counterfactual text in a plug-and-play fashion. In particular, given a text $\mathbf{x}$ and a control attribute $\mathbf{a}$, we seek to generate a controlled counterfactual $\mathbf{y}$. That is, we seek to draw samples from a distribution $p(\mathbf{y}|\mathbf{x}, \mathbf{a})$ where the generated sample $\mathbf{y}$ depends both on the input text $\mathbf{x}$ and the given control attribute $\mathbf{a}$. Hence, our task is different and more challenging than the simple controlled text generation task where the generated samples need to depend only on the control attribute $\mathbf{a}$.

Our main idea is as follows: Similar to how PPLM [9] steers a pretrained text generator $p(\mathbf{y}) \rightarrow p(\mathbf{y}|\mathbf{a})$ using the control attribute $\mathbf{a}$, we shall steer a pretrained text-to-text generator $p(\mathbf{y}|\mathbf{x}) \rightarrow p(\mathbf{y}|\mathbf{x}, \mathbf{a})$. In PPLM [9], the base model $p(\mathbf{y})$ is a pretrained GPT-2, while in our model, the base model $p(\mathbf{y}|\mathbf{x})$ is a pretrained BART model.

A BART model is a text-to-text model that takes as input a text $\mathbf{x}$ and produces a reconstruction $\mathbf{y}$ of the input text. The BART text-to-text framework consists of two modules: A BERT encoder and a GPT-2 decoder. That is, the model takes an input text and the BERT encoder first returns a text representation $\mathbf{e}$. This text representation is then given to the GPT-2 decoder to reconstruct the input text word-by-word. This can be summarized as follows:

$$\begin{aligned} \mathbf{e} &= \text{BERT}(\mathbf{x}), \\ H_t &= \text{Transformer}(\mathbf{y}_{<t}, \mathbf{e}), \\ \mathbf{o}_t &= \text{PredictionHead}(H_t), \\ \mathbf{y}_t &\sim \text{Categorical}(\mathbf{o}_t). \end{aligned}$$

where t is the word position in the text, H_t is the last hidden layer before the prediction head and $\mathbf{o}_t$ are the log-probabilities of the words in the vocabulary used for sampling the next word $\mathbf{y}_t$. We will refer to this as the unmodified BART model having the next word prediction distribution $p(\mathbf{y}_t|\mathbf{y}_{<t}, \mathbf{x})$. To steer the BART model, similarly to PPLM, we add a learnable perturbation ΔH_t to the hidden states H_t of the unmodified BART model. This can be summarized as follows:

$$\begin{aligned} \mathbf{e} &= \text{BERT}(\mathbf{x}), \\ H_t &= \text{Transformer}(\mathbf{y}_{<t}, \mathbf{e}), \\ \mathbf{o}_t &= \text{PredictionHead}(H_t + \Delta H_t), \\ \mathbf{y}_t &\sim \text{Categorical}(\mathbf{o}_t). \end{aligned}$$

We will refer to this as the modified BART model having the next word prediction distribution $\bar{p}(\mathbf{y}_t|\mathbf{y}_{<t}, \mathbf{x})$ Similarly to PPLM, the learning of the perturbation parameters $\{\Delta H_1, \ldots, \Delta H_T\}$ is driven by the following objective:

$$\begin{aligned} \mathcal{L}_{\text{CASPer}} = &-\log p(\mathbf{a}|\mathbf{y}) \\ &- \sum_{t=1}^{T} D_{\text{KL}}(p(\mathbf{y}_t|\mathbf{y}_{<t}, \mathbf{x})||\bar{p}(\mathbf{y}_t|\mathbf{y}_{<t}, \mathbf{x})). \end{aligned}$$

where the first term provides the learning signal to steer the counterfactuals towards the desired goal attribute by trying to maximize the log-probability of the desired attribute. The second term tries to keep the generations close to the unmodified BART to ensure that the text remains similar in content to the original input text and also remains fluent and plausible. We note that this learning process is done separately each time we need to generate a new sample. However, the learning of the perturbation parameters $\{\Delta H_1, \ldots, \Delta H_T\}$ can be done very quickly and it only takes about 100 gradient steps. This property makes CASPer a flexible way to generate counterfactuals of a given text.

Discussion. Note that if we simply obtain samples from a pretrained BART model $p(\mathbf{y}|\mathbf{x})$, a sample $\mathbf{y}$ can be considered as a counterfactual of the input text $\mathbf{x}$. However, this sample would

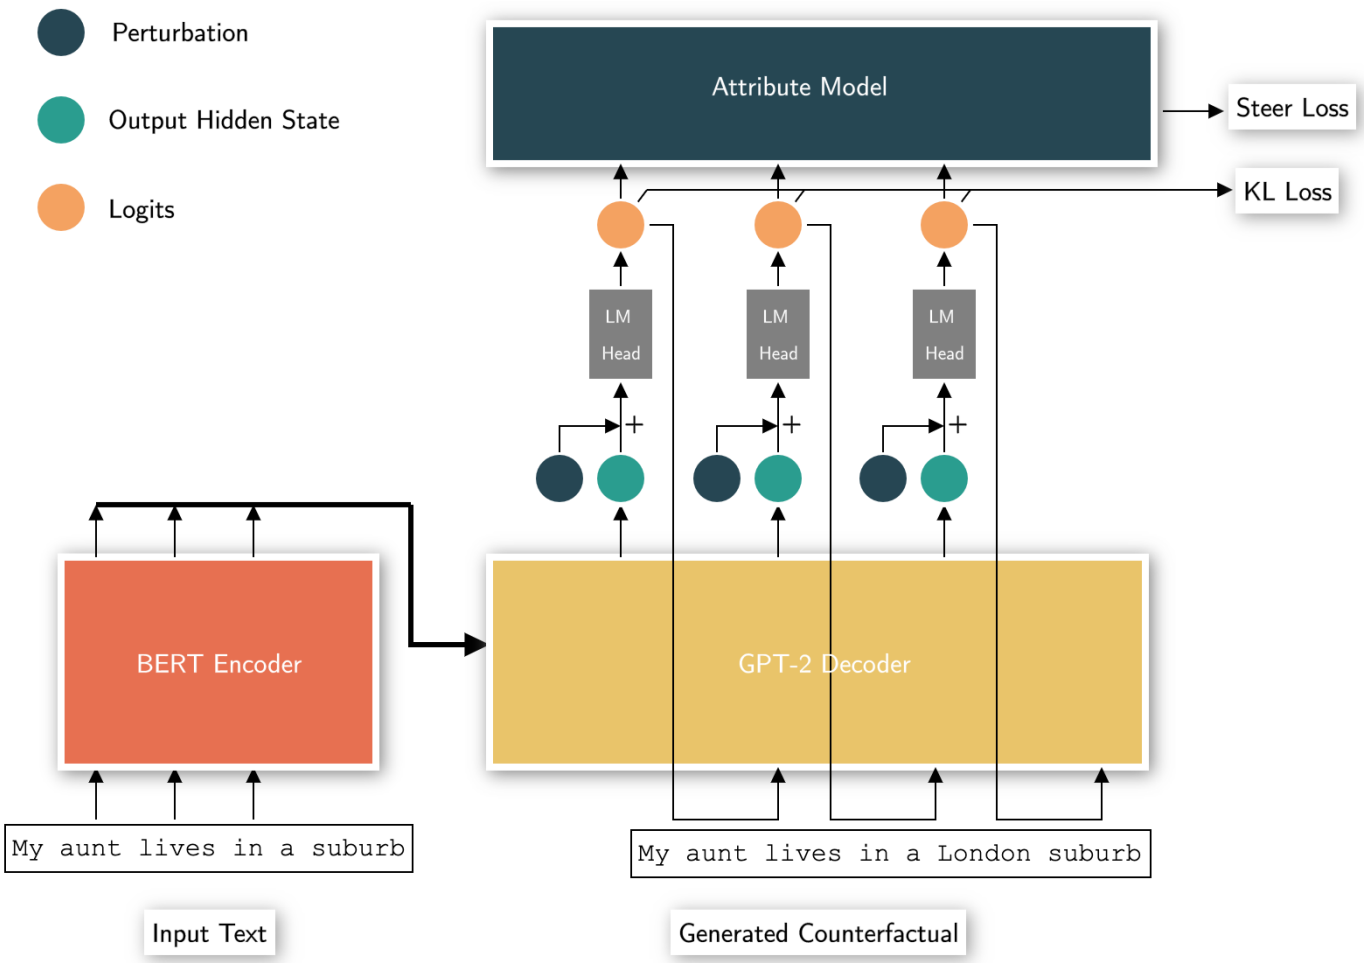

Fig. 1. Model Architecture of CASPer. The BERT encoder (shown in orange) takes the input text and returns a representation of the input text as a set of vectors. These are provided to the GPT-2 decoder for cross-attention. At each step of decoding, the decoder returns an output hidden state (shown as green circle). To this, we add a perturbation matrix (shown as blue circle). The perturbed hidden state is then provided to the language model head to obtain the logits over the vocabulary for sampling the next token. These logits are also provided to the attribute model for computing the loss for steering the generation. The KL divergence between the perturbed logits and the unperturbed logits is also minimized to keep the semantic content of the generated text close to the original input text.

be an almost exact reconstruction of the input text. Hence from the perspective of model testing, this type of counterfactual would not be much useful. However, by applying steering to $p(\mathbf{y}|\mathbf{x})$ using a control attribute $\mathbf{a}$, we are able to control in what way we want to modify the input text to generate the counterfactual. Hence, from the perspective of model testing, this type of counterfactual would be useful because we can test *how a deployed model will behave if the distribution of its inputs are perturbed in a certain way*.

IV. Related Work

The task of controlled text generation is well studied in literature. [17] propose a model aims to generate plausible sentences conditioned on representation vectors with semantic structure. Another work [16] focuses on controlled text generation, however, unlike the previous work [17], the conditioning need not be simply a class label. The conditioning can be a data structure such as a table. The model is trained end-to-end similarly to the objective of [17]. PPLM [9] combine a pre-trained language model, similarly to [1] with an attribute classifier to perform controlled language generation and use the attribute classifier to steer the text generation process without further training of any of the two models. [3], adopting a similar direction, deal with story completion with a desired sentiment. [37] is a model that controls text generation via 50 rigid control codes predetermined at training time. However all these works, cannot be used for counterfactual text generation as these are purely class-conditional generative models and do not allow generation conditioned on a given input text. Some earlier works, including and not limited to, [2], [9], [257]–[260] propose the idea of steering Language Models but these also can not be directly used for counterfactual generation task.

To tackle text-to-text generation tasks dealing with transfer of style or content, models such as [4]–[6] have been proposed. However, these works are not plug-and-play and lack the use of attribute model that can plugged flexibly at sampling time. This task of generating controlled counterfactuals has been attempted

in prior works by relying on template-matching and token-based substitutions to generate the test-cases [15], [26]. However, this can require significant human-involvement to curate the templates and the dictionaries. Hence, it cannot scale well when template and dictionaries need to be updated frequently. The work [15] employs a tool Checklist which is one of the attempts to come up with generalized perturbations. For generation, Checklist uses a set of pre-defined templates, lexicons, general-purpose perturbations, and context-aware suggestions. To better evaluate the deployed models, some prior works have relied on human designed test examples or either using templates [15], [155], [169], [190], [200], [238], [250]. Polyjuice [26], while seeking to automate the process, still requires paired dataset in the form of text and their perturbed versions for different control codes. Therefore the mapping between text and perturbed version is learned through supervision. Another parallel work Tailor [118] generates perturbations designed for different control codes by making use of a combination of semantic roles and content keywords. And thereby require supervision for different controls. In contrast, CASPer does not require any task-specific or control-code specific training and can be used to work with different control code models given input text. One work related to ours has been tackled in [107] which generates text samples given a text with a controlling that specify the scope of the generated text. LEWIS attempts to generate text perturbations by introducing blanks via template matching and filling in using pretrained language models [206]. However, this relies on rule-based template matching and human supervision to develop such templates. CATGen [116] tries to generate attribute-specific text but it requires training of sequence to sequence model with pre-determined control codes for perturbation. Hence, it lacks the flexibility of a plug-and-play approach like ours. MiCE [103] proposes a technique to generate counterfactual explanations which are human interpretable and user-centric. It fine-tunes a T5 model to generate counterfactual text and use them as explanations for the behavior of the deployed models but lack feature-attributions. Another work [118] tries to generate perturbation with semantic controls but rely on specific templates derived using semantic roles and other labeling heuristics. A work close to ours, GYC, the inference of latent representation of the input text with respect to a GPT-2 decoder is done directly via optimization. This approach fails to achieve good inference for long and complex text [107]. To target model failure, thus implicitly acting as a form of model testing, prior works have attempted the use of adversarial approaches [181], [198], [229] stemming from the need to build robust models via adversarial testing [42]–[44], [54]. However, these are still limited to specific deomains and generations are likely to be not plausible to be seen in the input text [55] or may require additional human effort [56]. Some works have attempted to change style attributes automatically either with no control or with predefined style templates [205], [208]. The notion of counterfactuals [8], [12] and their use in model testing for has also been applied towards testing in models that consume structured inputs [65], [66], [68]

V. Experiments

The goal of the experiments is to: *1)* show that a flexible plug-and-play framework can effectively achieve controlled counterfactual text generation and the generated text is fluent, plausible, diverse and follows the steering provided by the attribute model. *2)* We evaluate how well the generated perturbations can act as data-augmentation samples in order to make a downstream classification task performance more robust.

A. Datasets

We evaluate the models on the following data sets text.

1) **YELP Sentiment Dataset.** To evaluate how our model is able to change the sentiment of the original text and achieve the target sentiment, we use the YELP sentiment dataset [46]. This dataset is also characterized by informal text which can be seen in realistic user inputs.
2) **IMDB Sentiment Dataset.** To further evaluate how our model is able to change the sentiment of the original input text, we test on IMDB Sentiment Dataset [204]. This dataset is also characterized by long and complex text and is thus a challenging dataset.

B. Controlled Text Generation with Attribute Steering

We first evaluate the quality of our generated text with respect to the steering signal. We expect our generated text to preserve the semantic content and syntactic structure of the input text while being fluent and diverse as we steer the text towards the target attribute. The steering signal we evaluate in this work is to make the sentiment of the target text from negative to positive.

1) Baselines: To compare with state-of-the-art template-based methods relying on token substitutions via dictionaries we compare with Checklist [15]. We specifically consider the perturbation helper that relies on RoBERTa [49] to fill-in-the-blank. In comparison to this baseline, we expect ours to generate more fluent and diverse text samples that is free from the restrictions of the pre-specified templates. We also compare against Masked-LM [36], which is a dictionary-free approach but still relies on masking a specific token in the input text and and letting the model fill-in the masked token. The randomness in this filling-in process leads to generation of counterfactual samples. Hence this model generates only token-level substitutions and does not generate fluent sentence-level text. We expect CASPer to address this limitation of purely token-level substitutions.

To compare with state-of-the-art adversarial methods we compare with Hotflip [44] [42]. In comparison to this baseline, we expect ours to generate more fluent samples. We expect to see that content preservation of such approaches is high as these methods rely on changing the highest gradient word with another word that would flip the label.

To compare with state-of-the-art text generation methods we compare with PPLM [9] based on GPT-2 [19] and Polyjuice [26] based on finetuned GPT-2 [19]. In comparison to this baseline, we expect ours to generate more fluent and diverse

Metrics	Dataset	RoBERTa	Token-based	GPT−2	Gradient-based	Finetuning	Ours
		Mask−LM	Checklist	PPLM	Hotflip	Polyjuice	CASPer
CP ↑	YELP	0.30	0.321	0.064	0.365	0.212	0.202
	IMDB	0.29	0.30	0.048	0.291	0.317	0.231
Perplexity ↓	YELP	3.82	3.79	3.544	3.95	3.64	**3.44**
	IMDB	3.05	3.12	3.35	3.69	3.331	**2.80**
BLEU−4 ↓	YELP	0.903	0.530	0.064	NA	0.521	**0.309**
	IMDB	0.9027	0.909	0.042	NA	0.861	**0.231**

TABLE II
COMPARISON BETWEEN MODELS ON THE YELP AND IMDB DATASET. THE MODEL USED FOR STEERING IS A PRE-TRAINED SENTIMENT CLASSIFICATION MODEL.

samples. While for PPLM, since it simply takes a prompt text and completes the text, it has no incentive to generate text that preserves its content. Thus, we expect ours to preserve content better. In diversity of the samples, PPLM, because it is not tasked with preserving content, can generate arbitrary and overly diverse samples. Thus we highlight that while we perform a comparison of sample diversity between PPLM and our model, still the performances are not directly comparable because the expectations are different from both models. For Polyjuice, we expect ours to generate more fluent and diverse samples because, unlike Polyjuice, ours is free from specific template-based generation.

2) Metrics: To assess the quality of generated counterfactual text we focus on evaluating content preservation, fluency, diversity and syntactic similarity. We use the following metrics to measure the above characteristics.

1) **Content Preservation.** By measuring content preservation, we assess the similarity between input text and the counterfactual text samples. For this, we use the transformer model proposed in [45]. While higher content preservation is desirable in general, this metric alone does not provide the complete evaluation. Therefore for proper evaluation, we will introduce a second metric that measures sample diversity.
2) **Diversity.** This metric evaluates how different are the generated samples from each other. We find the BLEU-4 score between the input text and the generated text. Hence, if this score is lower, then the generated counterfactual samples have a high diversity at the token level.
3) **Fluency.** Fluency of the generated samples is important to evaluate because the samples must come from a distribution that the test model is likely to see when it is deployed. This is computed by finding the perplexity score of the generated output. We take a GPT-2 model for computing the perplexity. Lower perplexity implies that the generated text is more fluent.

3) Quantitative Results: In Tab. II, our results on perplexity show that CASPer outperforms the baselines significantly and achieves a lower perplexity score. This shows that the samples generated by CASPer are fluent and plausible. We also show that CASPer is able effectively preserve the content in its samples. We note that our model is competitive with rule-based token substitution methods like Checklist. Lastly, in terms of BLEU-4 score, we note that our model again outperforms the baselines with our generated samples achieving the lowest values of BLEU-4 score with the original input text only exception being PPLM. This shows that our model indeed generates diverse samples with that have low token-level match with the input text. In diversity of the samples, the baseline model PPLM, because it is not tasked with preserving content, can generate arbitrary and overly diverse samples. Thus we highlight that while we perform a comparison of sample diversity between PPLM and our model, still the performances are not directly comparable because the expectations are different from both models. Considering the performance on all the metrics together shows that CASPer is able to effectively generate samples that preserve the original content, are fluent and diverse in comparison to the baselines.

4) Qualitative Analysis: In Tab. III, we show samples of text generated by CASPer. We show two experiments. In the first experiment, our steering goal was to take an input text and perturb it so that the probability of its sentiment becomes large. The probability of the sentiment is estimated using a pre-trained sentiment classification model. The initial label of the text was negative. On the right we note that CASPer has successfully perturbed the text to change its sentiment label to positive. Furthermore, note that the generated sample have good content preservation as all the samples talk about movies and actors. Furthermore, the model makes some important changes to the content that result in a change in the sentiment of the text. We also note that each sample is different from the other thus producing diverse samples. Lastly, we note that the samples are fluent and plausible text samples.

In the second experiment, our steering goal was to take a sentence that does not contain a location named-entity and perturb it so that contains a location named-entity. We see that CASPer produces samples that contain a location named-entity tag. We also note that the named entity that the model introduces are diverse and are used in a variety of contexts in the generated text. As before, the text samples are fluent and preserve the content of the original input text. For this task, note that these samples clearly retain the sentiment of the text and only introduce some location entities. Because we expect the actual location (i.e. UK or Libya) should not be a causal term in prediction of the sentiment of the text, these samples can act as effective samples for augmenting the training data when we train a downstream sentiment model. While a model that is

Inputs	Generated Controlled Perturbations
Steering Goal : To maximize the probability of the positive class label with respect to a pre-trained sentiment classifier. **Input Text**: `We know from other movies that the writers can't solve problems, like this one. The actors are also good.` **Initial Label** : Negative	We know from other movies that the actors are good but they cannot help the movie. The premise was not too bad and is enjoyable. **New Label** : Positive We know from other documentaries that the acting was quite strong. We can't say the premise was good but the execution was good. **New Label** : Positive We know from other movies that the actors are good but they must win over the fans. The premise was not too bad. **New Label** : Positive We know from other movies that the actors are good. They can not help the movie. The premise was not too bad. **New Label** : Positive
Steering Goal: To make the sentence contain at least one location named-entity. **Input Text**: `A wonderful little production.The filming technique is very unassuming- very old-time-BBC fashion.` **Initial State** : No location named-entity is present in the text	A wonderful first look at this film from the UK. The filming is unassertive. ¡/s¿io is also one of the best productions. **Location Named-Entity**: UK A wonderful little Theatre is an old-time BBC production, set in an old London apartment block. The film is set in London landmarks **Location Named-Entity**: London A Wonderful Little Edition by BBC Films - set, out in UK now via iPlayer. A wonderful little production with some nice twists. **Location Named-Entity**: UK A wonderful little series about a BBC crew in Africa. Well-executed by the team who are experienced in film and rigging procedures. **Location Named-Entity**: Africa A wonderful little show about BBC's wartime bombing of Libya contributed by Peter Bradshaw in 1959. **Location Named-Entity**: Libya
Steering Goal : To maximize the probability of the positive class label with respect to a pre-trained sentiment classifier. **Input Text**: `I had the terrible misfortune of having to view this "b-movie" in it's entirety. All I have to say is--- save your time and money!!` **Initial Label** : Negative	I have the wonderful misfortune of having to view this New Girl in it's entirety. To view it positively, I must open up and say. **New Label** : Positive I had the terrible misfortune of having to view this transform into... in it's entirety. It's truly amazing how some programs remain.. **New Label** : Positive I had the terrible misfortune of having to view this movie in it endearing ways. It's a wonderful salute wives lenders acknowledge.. **New Label** : Positive I had the terrible misfortune of having to view this in it's full glory. All I have to say is wow. everyone agrees. **New Label** : Positive
Steering Goal : To maximize the probability of the negative class label with respect to a pre-trained sentiment classifier. **Input Text**: `For those who'd like to see this movie? I'd say: go! Without the narration it might be a very good movie/documentary.` **Initial Label** : Positive	For those who'd like to see it I suggest that they see this movie without narration. The movie isn't a very good movie. **New Label** : Negative For those want to get to know the movie through its narration? I'd say: go! Without the narration it might be a very bad movie. **New Label** : Negative

TABLE III
GENERATED CONTROLLED PERTURBATIONS FROM THE PROPOSED MODEL CASPER.

Model	Dataset	CASPer
Accuracy - No Aug	YELP	89.90
	IMDB	90.10
Accuracy - With Aug	YELP	**92.00**
	IMDB	**91.20**

TABLE IV
COMPARISON OF ACCURACY BETWEEN MODELS ON THE YELP AND IMDB DATASET. THE GENERATED DATA FOR NER TASK ON STEERING IS USED FOR ROBUSTIFYING AN N-GRAM BASED SENTIMENT MODEL.

biased may predict different labels based on the actual location token used, this kind of data augmentation will regularize the model to be more robust to such changes which should ideally not affect the predicted label of the test model.

C. Controlled Text for Model Robustification

In this section, we evaluate how well our generated samples can improve robustness of the test classifier. For this, we generated text samples to introduce a location named-entity in the input text. We assume that simply introducing a location named-entity should not change the class label of the text with respect to the test model. Hence, after generating the controlled perturbations, we take the original label of the input text from the training set and assign the same label to the generated samples. These new examples are added to the training set and producing data-augmented training set. Using this augmented training set, we then train the test model.

1) Baselines and Metrics: We generate samples using CASPer. We augment the generated samples to the training set and train the test model. We compared the accuracy of the test model trained without data augmentation and then trained with data augmentation via our counterfactual generation method.

2) Quantitative Results: In Table IV we show a comparison between the models. We note that the samples generated by CASPer using NER model are effective in robustifying the test model and produces significant improvement in the accuracy as compared to when training with original samples.

VI. CONCLUSION

In this paper, we introduced CASPer, a plug-and-play counterfactual text generation framework. We showed that our generated controlled perturbations preserve the content of the original text while also being fluent, diverse and effective in terms of the provided steering signal flexibly. We showed that samples generated by CASPer can act as effective candidates for training data augmentation and improve the robustness of the target model and preventing the target model from modeling spurious correlations between the target label and non-causal aspects of the input text.

REFERENCES

[1] Nguyen, A., Clune, J., Bengio, Y., Dosovitskiy, A. & Yosinski, J. Plug play generative networks: Conditional iterative generation of images in latent space. *Proceedings Of The IEEE Conference On Computer Vision And Pattern Recognition*. pp. 4467-4477 (2017)

[2] Krause, B., Gotmare, A., McCann, B., Keskar, N., Joty, S., Socher, R. & Rajani, N. Gedi: Generative discriminator guided sequence generation. *ArXiv Preprint ArXiv:2009.06367*. (2020)

[3] Luo, F., Dai, D., Yang, P., Liu, T., Chang, B., Sui, Z. & Sun, X. Learning to Control the Fine-grained Sentiment for Story Ending Generation. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6020-6026 (2019,7), https://aclanthology.org/P19-1603

[4] Lample, G., Subramanian, S., Smith, E., Denoyer, L., Ranzato, M. & Boureau, Y. Multiple-attribute text rewriting. *International Conference On Learning Representations*. (2018)

[5] Li, J., Jia, R., He, H. & Liang, P. Delete, retrieve, generate: A simple approach to sentiment and style transfer. *ArXiv Preprint ArXiv:1804.06437*. (2018)

[6] Shen, T., Lei, T., Barzilay, R. & Jaakkola, T. Style transfer from non-parallel text by cross-alignment. *ArXiv Preprint ArXiv:1705.09655*. (2017)

[7] Gusfield, D. Algorithms on Strings, Trees and Sequences. (Cambridge University Press,1997)

[8] Mothilal, R., Sharma, A. & Tan, C. Explaining machine learning classifiers through diverse counterfactual explanations. *Proceedings Of The 2020 Conference On Fairness, Accountability, And Transparency*. pp. 607-617 (2020)

[9] Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J. & Liu, R. Plug and play language models: a simple approach to controlled text generation. *ArXiv Preprint ArXiv:1912.02164*. (2019)

[10] Garg, S., Perot, V., Limtiaco, N., Taly, A., Chi, E. & Beutel, A. Counterfactual fairness in text classification through robustness. *Proceedings Of The 2019 AAAI/ACM Conference On AI, Ethics, And Society*. pp. 219-226 (2019)

[11] Huang, P., Zhang, H., Jiang, R., Stanforth, R., Welbl, J., Rae, J., Maini, V., Yogatama, D. & Kohli, P. Reducing Sentiment Bias in Language Models via Counterfactual Evaluation. *ArXiv Preprint ArXiv:1911.03064*. (2019)

[12] Wachter, S., Mittelstadt, B. & Russell, C. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *Harv. JL Tech.*. **31** pp. 841 (2017)

[13] Gusfield, D. Algorithms on Strings, Trees and Sequences. (Cambridge University Press,1997)

[14] Lee, K., Chang, M. & Toutanova, K. Latent Retrieval for Weakly Supervised Open Domain Question Answering. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6086-6096 (2019,7), https://www.aclweb.org/anthology/P19-1612

[15] Ribeiro, M., Wu, T., Guestrin, C. & Singh, S. Beyond accuracy: Behavioral testing of NLP models with CheckList. *ArXiv Preprint ArXiv:2005.04118*. (2020)

[16] Ye, R., Shi, W., Zhou, H., Wei, Z. & Li, L. Variational Template Machine for Data-to-Text Generation. *ArXiv Preprint ArXiv:2002.01127*. (2020)

[17] Hu, Z., Yang, Z., Liang, X., Salakhutdinov, R. & Xing, E. Toward controlled generation of text. *Proceedings Of The 34th International Conference On Machine Learning-Volume 70*. pp. 1587-1596 (2017)

[18] Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K. & Zettlemoyer, L. Deep Contextualized Word Representations. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. (2018), http://dx.doi.org/10.18653/v1/N18-1202

[19] Radford, A., Wu, J., Child, R., Luan, D., Amodei, D. & Sutskever, I. Language Models are Unsupervised Multitask Learners. *ArXiv*. (2019)

[20] Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M. & Hajishirzi, H. Text generation from knowledge graphs with graph transformers. *ArXiv Preprint ArXiv:1904.02342*. (2019)

[21] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł. & Polosukhin, I. Attention is All you Need. *Advances In Neural Information Processing Systems 30*. pp. 5998-6008 (2017), http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf

[22] Ficler, J. & Goldberg, Y. Controlling Linguistic Style Aspects in Neural Language Generation. *ArXiv*. **abs/1707.02633** (2017)

[23] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: sequence generative adversarial nets with policy gradient. (2017,8)

[24] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-Tuning Language Models from Human Preferences. (2019,9)

[25] Subramani, N., Bowman, S. & Cho, K. Can Unconditional Language Models Recover Arbitrary Sentences?. *NeurIPS*. (2019)

[26] Wu, T., Ribeiro, M., Heer, J. & Weld, D. Polyjuice: Automated, General-purpose Counterfactual Generation. *ArXiv Preprint ArXiv:2101.00288*. (2021)

[27] Huang, P., Zhang, H., Jiang, R., Stanforth, R., Welbl, J., Rae, J., Maini, V., Yogatama, D. & Kohli, P. Reducing Sentiment Bias in Language Models via Counterfactual Evaluation. *ArXiv Preprint ArXiv:1911.03064*. (2019)

[28] Lee, K., Chang, M. & Toutanova, K. Latent Retrieval for Weakly Supervised Open Domain Question Answering. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6086-6096 (2019,7), https://www.aclweb.org/anthology/P19-1612

[29] Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K. & Zettlemoyer, L. Deep Contextualized Word Representations. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. (2018), http://dx.doi.org/10.18653/v1/N18-1202

[30] Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M. & Hajishirzi, H. Text generation from knowledge graphs with graph transformers. *ArXiv Preprint ArXiv:1904.02342*. (2019)

[31] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł. & Polosukhin, I. Attention is All you Need. *Advances In Neural Information Processing Systems 30*. pp. 5998-6008 (2017), http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf

[32] Ficler, J. & Goldberg, Y. Controlling Linguistic Style Aspects in Neural Language Generation. *ArXiv*. **abs/1707.02633** (2017)

[33] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: sequence generative adversarial nets with policy gradient. (2017,8)

[34] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-Tuning Language Models from Human Preferences. (2019,9)

[35] Subramani, N., Bowman, S. & Cho, K. Can Unconditional Language Models Recover Arbitrary Sentences?. *NeurIPS*. (2019)

[36] Devlin, J., Chang, M., Lee, K. & Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *ArXiv Preprint ArXiv:1810.04805*. (2018)

[37] Keskar, N., McCann, B., Varshney, L., Xiong, C. & Socher, R. Ctrl: A conditional transformer language model for controllable generation. *ArXiv Preprint ArXiv:1909.05858*. (2019)

[38] Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *ArXiv:1301.3781*. (2013)

[39] Howard, J. & Ruder, S. Universal language model fine-tuning for text classification. *ArXiv Preprint ArXiv:1801.06146*. (2018)

[40] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-tuning language models from human preferences. *ArXiv Preprint ArXiv:1909.08593*. (2019)

[41] Ribeiro, M., Singh, S. & Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations.. *AAAI*. **18** pp. 1527-1535 (2018)

[42] Michel, P., Li, X., Neubig, G. & Pino, J. On evaluation of adversarial perturbations for sequence-to-sequence models. *ArXiv Preprint ArXiv:1903.06620*. (2019)

[43] Zhao, Z., Dua, D. & Singh, S. Generating natural adversarial examples. *ArXiv Preprint ArXiv:1710.11342*. (2017)

[44] Ebrahimi, J., Rao, A., Lowd, D. & Dou, D. Hotflip: White-box adversarial examples for text classification. *ArXiv Preprint ArXiv:1712.06751*. (2017)

[45] Reimers, N. & Gurevych, I. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing*. (2019,11), http://arxiv.org/abs/1908.10084

[46] Zhang, X., Zhao, J. & LeCun, Y. Character-level convolutional networks for text classification. *Advances In Neural Information Processing Systems*. pp. 649-657 (2015)

[47] Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P., Hellmann, S., Morsey, M., Van Kleef, P., Auer, S. & Bizer, C. DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic Web Journal*. **6** (2014,1)

[48] Gardner, M., Grus, J., Neumann, M., Tafjord, O., Dasigi, P., Liu, N., Peters, M., Schmitz, M. & Zettlemoyer, L. AllenNLP: A deep semantic natural language processing platform, 2018. *ArXiv Preprint ArXiv:1803.07640*. (2017)

[49] Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L. & Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *ArXiv Preprint ArXiv:1907.11692*. (2019)

[50] Zhang, T., Kishore, V., Wu, F., Weinberger, K. & Artzi, Y. BERTScore: Evaluating Text Generation with BERT. *International Conference On Learning Representations*. (2020), https://openreview.net/forum?id=SkeHuCVFDr

[51] Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. & Le, Q. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances In Neural Information Processing Systems*. pp. 5753-5763 (2019)

[52] Ma, P., Wang, S. & Liu, J. Metamorphic Testing and Certified Mitigation of Fairness Violations in NLP Models. *Proceedings Of The Twenty-Ninth International Joint Conference On Artificial Intelligence, IJCAI-20*. pp. 458-465 (2020,7), Main track

[53] Holstein, K., Wortman Vaughan, J., Daumé III, H., Dudik, M. & Wallach, H. Improving fairness in machine learning systems: What do industry practitioners need?. *Proceedings Of The 2019 CHI Conference On Human Factors In Computing Systems*. pp. 1-16 (2019)

[54] Goodfellow, I., Shlens, J. & Szegedy, C. Explaining and harnessing adversarial examples. *ArXiv Preprint ArXiv:1412.6572*. (2014)

[55] Li, J., Monroe, W. & Jurafsky, D. Understanding neural networks through representation erasure. *ArXiv Preprint ArXiv:1612.08220*. (2016)

[56] Jia, R. & Liang, P. Adversarial examples for evaluating reading comprehension systems. *ArXiv Preprint ArXiv:1707.07328*. (2017)

[57] Jang, E., Gu, S. & Poole, B. Categorical reparameterization with gumbel-softmax. *ArXiv Preprint ArXiv:1611.01144*. (2016)

[58] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. *Thirty-first AAAI Conference On Artificial Intelligence*. (2017)

[59] Zhang, Y., Gan, Z., Fan, K., Chen, Z., Henao, R., Shen, D. & Carin, L. Adversarial feature matching for text generation. *ArXiv Preprint ArXiv:1706.03850*. (2017)

[60] Engel, J., Hoffman, M. & Roberts, A. Latent constraints: Learning to generate conditionally from unconditional generative models. *ArXiv Preprint ArXiv:1711.05772*. (2017)

[61] Zhang, K. & Shasha, D. Simple fast algorithms for the editing distance between trees and related problems. *SIAM Journal On Computing*. **18**, 1245-1262 (1989)

[62] Pearl, J. & Others Models, reasoning and inference. *Cambridge, UK: CambridgeUniversityPress*. (2000)

[63] Madaan, N., Mehta, S., Agrawaal, T., Malhotra, V., Aggarwal, A., Gupta, Y. & Saxena, M. Analyze, detect and remove gender stereotyping from bollywood movies. *Conference On Fairness, Accountability And Transparency*. pp. 92-105 (2018)

[64] Aggarwal, A., Lohia, P., Nagar, S., Dey, K. & Saha, D. Black box fairness testing of machine learning models. *Proceedings Of The 2019 27th ACM Joint Meeting On European Software Engineering Conference And Symposium On The Foundations Of Software Engineering*. pp. 625-635 (2019)

[65] Galhotra, S., Brun, Y. & Meliou, A. Fairness testing: testing software for discrimination. *Proceedings Of The 2017 11th Joint Meeting On Foundations Of Software Engineering*. pp. 498-510 (2017)

[66] Udeshi, S., Arora, P. & Chattopadhyay, S. Automated directed fairness testing. *Proceedings Of The 33rd ACM/IEEE International Conference On Automated Software Engineering*. pp. 98-108 (2018)

[67] Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J. & Yu, Y. Texygen: A benchmarking platform for text generation models. *The 41st International ACM SIGIR Conference On Research Development In Information Retrieval*. pp. 1097-1100 (2018)

[68] John, P., Vijaykeerthy, D. & Saha, D. Verifying Individual Fairness in Machine Learning Models. *Conference On Uncertainty In Artificial Intelligence*. pp. 749-758 (2020)

[69] Madaan, N., Singh, G., Mehta, S., Chetan, A. & Joshi, B. Generating clues for gender based occupation de-biasing in text. *ArXiv Preprint ArXiv:1804.03839*. (2018)

[70] Madaan, N., Mehta, S., Agrawaal, T., Malhotra, V., Aggarwal, A. & Saxena, M. Analyzing Gender Stereotyping in Bollywood Movies. *ArXiv Preprint ArXiv:1710.04117*. (2017)

[71] Suvarna, A., Dey, K., Nagar, S., Madaan, N. & Mehta, S. Handling Gender Biases in E-Commerce Product Specifications. *2019 Grace Hopper Celebration India (GHCI)*. pp. 1-5 (2019)

[72] Madaan, N., Mehta, S., Mittal, S. & Suvarna, A. Judging a Book by its Description: Analyzing Gender Stereotypes in the Man Bookers Prize Winning Fiction. *ArXiv Preprint ArXiv:1807.10615*. (2018)

[73] Huang, P., Zhang, H., Jiang, R., Stanforth, R., Welbl, J., Rae, J., Maini, V., Yogatama, D. & Kohli, P. Reducing Sentiment Bias in Language Models via Counterfactual Evaluation. *ArXiv Preprint ArXiv:1911.03064*. (2019)

[74] Lee, K., Chang, M. & Toutanova, K. Latent Retrieval for Weakly Supervised Open Domain Question Answering. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6086-6096 (2019,7), https://www.aclweb.org/anthology/P19-1612

[75] Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K. & Zettlemoyer, L. Deep Contextualized Word Representations. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. (2018), http://dx.doi.org/10.18653/v1/N18-1202

[76] Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M. & Hajishirzi, H. Text generation from knowledge graphs with graph transformers. *ArXiv Preprint ArXiv:1904.02342*. (2019)

[77] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł. & Polosukhin, I. Attention is All you Need. *Advances In Neural Information Processing Systems 30*. pp. 5998-6008 (2017), http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf

[78] Ficler, J. & Goldberg, Y. Controlling Linguistic Style Aspects in Neural Language Generation. *ArXiv*. **abs/1707.02633** (2017)

[79] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: sequence generative adversarial nets with policy gradient. (2017,8)

[80] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-Tuning Language Models from Human Preferences. (2019,9)

[81] Subramani, N., Bowman, S. & Cho, K. Can Unconditional Language Models Recover Arbitrary Sentences?. *NeurIPS*. (2019)

[82] Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *ArXiv:1301.3781*. (2013)

[83] Howard, J. & Ruder, S. Universal language model fine-tuning for text classification. *ArXiv Preprint ArXiv:1801.06146*. (2018)

[84] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-tuning language models from human preferences. *ArXiv Preprint ArXiv:1909.08593*. (2019)

[85] Ribeiro, M., Singh, S. & Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations.. *AAAI*. **18** pp. 1527-1535 (2018)

[86] Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P., Hellmann, S., Morsey, M., Van Kleef, P., Auer, S. & Bizer, C. DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic Web Journal*. **6** (2014,1)

[87] Gardner, M., Grus, J., Neumann, M., Tafjord, O., Dasigi, P., Liu, N., Peters, M., Schmitz, M. & Zettlemoyer, L. AllenNLP: A deep semantic natural language processing platform, 2018. *ArXiv Preprint ArXiv:1803.07640*. (2017)

[88] Zhang, T., Kishore, V., Wu, F., Weinberger, K. & Artzi, Y. BERTScore: Evaluating Text Generation with BERT. *International Conference On Learning Representations*. (2020), https://openreview.net/forum?id=SkeHuCVFDr

[89] Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. & Le, Q. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances In Neural Information Processing Systems*. pp. 5753-5763 (2019)

[90] Jang, E., Gu, S. & Poole, B. Categorical reparameterization with gumbel-softmax. *ArXiv Preprint ArXiv:1611.01144*. (2016)

[91] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. *Thirty-first AAAI Conference On Artificial Intelligence*. (2017)

[92] Zhang, Y., Gan, Z., Fan, K., Chen, Z., Henao, R., Shen, D. & Carin, L. Adversarial feature matching for text generation. *ArXiv Preprint ArXiv:1706.03850*. (2017)

[93] Engel, J., Hoffman, M. & Roberts, A. Latent constraints: Learning to generate conditionally from unconditional generative models. *ArXiv Preprint ArXiv:1711.05772*. (2017)

[94] Zhang, K. & Shasha, D. Simple fast algorithms for the editing distance between trees and related problems. *SIAM Journal On Computing*. **18**, 1245-1262 (1989)

[95] Madaan, N., Mehta, S., Agrawaal, T., Malhotra, V., Aggarwal, A., Gupta, Y. & Saxena, M. Analyze, detect and remove gender stereotyping from bollywood movies. *Conference On Fairness, Accountability And Transparency*. pp. 92-105 (2018)

[96] Aggarwal, A., Lohia, P., Nagar, S., Dey, K. & Saha, D. Black box fairness testing of machine learning models. *Proceedings Of The 2019 27th ACM Joint Meeting On European Software Engineering Conference And Symposium On The Foundations Of Software Engineering*. pp. 625-635 (2019)

[97] Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J. & Yu, Y. Texygen: A benchmarking platform for text generation models. *The 41st International ACM SIGIR Conference On Research Development In Information Retrieval*. pp. 1097-1100 (2018)

[98] Madaan, N., Singh, G., Mehta, S., Chetan, A. & Joshi, B. Generating clues for gender based occupation de-biasing in text. *ArXiv Preprint ArXiv:1804.03839*. (2018)

[99] Madaan, N., Mehta, S., Agrawaal, T., Malhotra, V., Aggarwal, A. & Saxena, M. Analyzing Gender Stereotyping in Bollywood Movies. *ArXiv Preprint ArXiv:1710.04117*. (2017)

[100] Suvarna, A., Dey, K., Nagar, S., Madaan, N. & Mehta, S. Handling Gender Biases in E-Commerce Product Specifications. *2019 Grace Hopper Celebration India (GHCI)*. pp. 1-5 (2019)

[101] Madaan, N., Mehta, S., Mittal, S. & Suvarna, A. Judging a Book by its Description: Analyzing Gender Stereotypes in the Man Bookers Prize Winning Fiction. *ArXiv Preprint ArXiv:1807.10615*. (2018)

[102] Gusfield, D. Algorithms on Strings, Trees and Sequences. (Cambridge University Press,1997)

[103] Ross, A., Marasović, A. & Peters, M. Explaining nlp models via minimal contrastive editing (mice). *ArXiv Preprint ArXiv:2012.13985*. (2020)

[104] Huang, P., Zhang, H., Jiang, R., Stanforth, R., Welbl, J., Rae, J., Maini, V., Yogatama, D. & Kohli, P. Reducing Sentiment Bias in Language Models via Counterfactual Evaluation. *ArXiv Preprint ArXiv:1911.03064*. (2019)

[105] Gusfield, D. Algorithms on Strings, Trees and Sequences. (Cambridge University Press,1997)

[106] Lee, K., Chang, M. & Toutanova, K. Latent Retrieval for Weakly Supervised Open Domain Question Answering. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6086-6096 (2019,7), https://www.aclweb.org/anthology/P19-1612

[107] Madaan, N., Padhi, I., Panwar, N. & Saha, D. Generate Your Counterfactuals: Towards Controlled Counterfactual Generation for Text. *Proceedings Of The AAAI Conference On Artificial Intelligence*., 13516-13524 (2021)

[108] Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K. & Zettlemoyer, L. Deep Contextualized Word Representations. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. (2018), http://dx.doi.org/10.18653/v1/N18-1202

[109] Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M. & Hajishirzi, H. Text generation from knowledge graphs with graph transformers. *ArXiv Preprint ArXiv:1904.02342*. (2019)

[110] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł. & Polosukhin, I. Attention is All you Need. *Advances In Neural Information Processing Systems 30*. pp. 5998-6008 (2017), http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf

[111] Ficler, J. & Goldberg, Y. Controlling Linguistic Style Aspects in Neural Language Generation. *ArXiv*. **abs/1707.02633** (2017)

[112] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: sequence generative adversarial nets with policy gradient. (2017,8)

[113] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-Tuning Language Models from Human Preferences. (2019,9)

[114] Subramani, N., Bowman, S. & Cho, K. Can Unconditional Language Models Recover Arbitrary Sentences?. *NeurIPS*. (2019)

[115] Huang, P., Zhang, H., Jiang, R., Stanforth, R., Welbl, J., Rae, J., Maini, V., Yogatama, D. & Kohli, P. Reducing Sentiment Bias in Language Models via Counterfactual Evaluation. *ArXiv Preprint ArXiv:1911.03064*. (2019)

[116] Wang, T., Wang, X., Qin, Y., Packer, B., Li, K., Chen, J., Beutel, A. & Chi, E. CAT-Gen: Improving Robustness in NLP Models via Controlled Adversarial Text Generation. *ArXiv Preprint ArXiv:2010.02338*. (2020)

[117] Radford, A., Narasimhan, K., Salimans, T. & Sutskever, I. Improving language understanding by generative pre-training (2018). (2018)

[118] Ross, A., Wu, T., Peng, H., Peters, M. & Gardner, M. Tailor: Generating and Perturbing Text with Semantic Controls. *ArXiv Preprint ArXiv:2107.07150*. (2021)

[119] Lee, K., Chang, M. & Toutanova, K. Latent Retrieval for Weakly Supervised Open Domain Question Answering. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6086-6096 (2019,7), https://www.aclweb.org/anthology/P19-1612

[120] Peters, M., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K. & Zettlemoyer, L. Deep Contextualized Word Representations. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. (2018), http://dx.doi.org/10.18653/v1/N18-1202

[121] Koncel-Kedziorski, R., Bekal, D., Luan, Y., Lapata, M. & Hajishirzi, H. Text generation from knowledge graphs with graph transformers. *ArXiv Preprint ArXiv:1904.02342*. (2019)

[122] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, Ł. & Polosukhin, I. Attention is All you Need. *Advances In Neural Information Processing Systems 30*. pp. 5998-6008 (2017), http://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf

[123] Ficler, J. & Goldberg, Y. Controlling Linguistic Style Aspects in Neural Language Generation. *ArXiv*. **abs/1707.02633** (2017)

[124] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: sequence generative adversarial nets with policy gradient. (2017,8)

[125] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-Tuning Language Models from Human Preferences. (2019,9)

[126] Subramani, N., Bowman, S. & Cho, K. Can Unconditional Language Models Recover Arbitrary Sentences?. *NeurIPS*. (2019)

[127] Mikolov, T., Chen, K., Corrado, G. & Dean, J. Efficient estimation of word representations in vector space. *ArXiv:1301.3781*. (2013)

[128] Howard, J. & Ruder, S. Universal language model fine-tuning for text classification. *ArXiv Preprint ArXiv:1801.06146*. (2018)

[129] Ziegler, D., Stiennon, N., Wu, J., Brown, T., Radford, A., Amodei, D., Christiano, P. & Irving, G. Fine-tuning language models from human preferences. *ArXiv Preprint ArXiv:1909.08593*. (2019)

[130] Ribeiro, M., Singh, S. & Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations.. *AAAI*. **18** pp. 1527-1535 (2018)

[131] Lehmann, J., Isele, R., Jakob, M., Jentzsch, A., Kontokostas, D., Mendes, P., Hellmann, S., Morsey, M., Van Kleef, P., Auer, S. & Bizer, C. DBpedia - A Large-scale, Multilingual Knowledge Base Extracted from Wikipedia. *Semantic Web Journal*. **6** (2014,1)

[132] Gardner, M., Grus, J., Neumann, M., Tafjord, O., Dasigi, P., Liu, N., Peters, M., Schmitz, M. & Zettlemoyer, L. AllenNLP: A deep semantic natural language processing platform, 2018. *ArXiv Preprint ArXiv:1803.07640*. (2017)

[133] Zhang, T., Kishore, V., Wu, F., Weinberger, K. & Artzi, Y. BERTScore: Evaluating Text Generation with BERT. *International Conference On Learning Representations*. (2020), https://openreview.net/forum?id=SkeHuCVFDr

[134] Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R. & Le, Q. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances In Neural Information Processing Systems*. pp. 5753-5763 (2019)

[135] Jang, E., Gu, S. & Poole, B. Categorical reparameterization with gumbel-softmax. *ArXiv Preprint ArXiv:1611.01144*. (2016)

[136] Yu, L., Zhang, W., Wang, J. & Yu, Y. Seqgan: Sequence generative adversarial nets with policy gradient. *Thirty-first AAAI Conference On Artificial Intelligence*. (2017)

[137] Zhang, Y., Gan, Z., Fan, K., Chen, Z., Henao, R., Shen, D. & Carin, L. Adversarial feature matching for text generation. *ArXiv Preprint ArXiv:1706.03850*. (2017)

[138] Engel, J., Hoffman, M. & Roberts, A. Latent constraints: Learning to generate conditionally from unconditional generative models. *ArXiv Preprint ArXiv:1711.05772*. (2017)

[139] Zhang, K. & Shasha, D. Simple fast algorithms for the editing distance between trees and related problems. *SIAM Journal On Computing*. **18**, 1245-1262 (1989)

[140] Madaan, N., Mehta, S., Agrawaal, T., Malhotra, V., Aggarwal, A., Gupta, Y. & Saxena, M. Analyze, detect and remove gender stereotyping from bollywood movies. *Conference On Fairness, Accountability And Transparency*. pp. 92-105 (2018)

[141] Aggarwal, A., Lohia, P., Nagar, S., Dey, K. & Saha, D. Black box fairness testing of machine learning models. *Proceedings Of The 2019 27th ACM Joint Meeting On European Software Engineering Conference And Symposium On The Foundations Of Software Engineering*. pp. 625-635 (2019)

[142] Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J. & Yu, Y. Texygen: A benchmarking platform for text generation models. *The 41st International ACM SIGIR Conference On Research Development In Information Retrieval*. pp. 1097-1100 (2018)

[143] Madaan, N., Singh, G., Mehta, S., Chetan, A. & Joshi, B. Generating clues for gender based occupation de-biasing in text. *ArXiv Preprint ArXiv:1804.03839*. (2018)

[144] Madaan, N., Mehta, S., Agrawaal, T., Malhotra, V., Aggarwal, A. & Saxena, M. Analyzing Gender Stereotyping in Bollywood Movies. *ArXiv Preprint ArXiv:1710.04117*. (2017)

[145] Suvarna, A., Dey, K., Nagar, S., Madaan, N. & Mehta, S. Handling Gender Biases in E-Commerce Product Specifications. *2019 Grace Hopper Celebration India (GHCI)*. pp. 1-5 (2019)

[146] Madaan, N., Mehta, S., Mittal, S. & Suvarna, A. Judging a Book by its Description: Analyzing Gender Stereotypes in the Man Bookers Prize Winning Fiction. *ArXiv Preprint ArXiv:1807.10615*. (2018)

[147] Beck, S., Riggs, K. & Gorniak, S. Relating developments in children's counterfactual thinking and executive functions. *Thinking Reasoning*. **15**, 337-354 (2009)

[148] Buchsbaum, D., Bridgers, S., Skolnick Weisberg, D. & Gopnik, A. The power of possibility: Causal learning, counterfactual reasoning, and pretend play. *Philosophical Transactions Of The Royal Society B: Biological Sciences*. **367**, 2202-2212 (2012)

[149] Weisberg, D. & Gopnik, A. Pretense, counterfactuals, and Bayesian causal models: Why what is not real really matters. *Cognitive Science*. **37**, 1368-1381 (2013)

[150] Dhurandhar, A., Chen, P., Luss, R., Tu, C., Ting, P., Shanmugam, K. & Das, P. Explanations based on the missing: Towards contrastive explanations with pertinent negatives. *ArXiv Preprint ArXiv:1802.07623*. (2018)

[151] Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V. & Zettlemoyer, L. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *ArXiv Preprint ArXiv:1910.13461*. (2019)

[152] Zhang, Y., Baldridge, J. & He, L. PAWS: Paraphrase adversaries from word scrambling. *ArXiv Preprint ArXiv:1904.01130*. (2019)

[153] Alzantot, M., Sharma, Y., Elgohary, A., Ho, B., Srivastava, M. & Chang, K. Generating Natural Language Adversarial Examples. *Proceedings Of The 2018 Conference On Empirical Methods In Natural Language Processing*. pp. 2890-2896 (2018), https://www.aclweb.org/anthology/D18-1316

[154] Alzantot, M., Sharma, Y., Elgohary, A., Ho, B., Srivastava, M. & Chang, K. Generating Natural Language Adversarial Examples. *Proceedings Of The 2018 Conference On Empirical Methods In Natural Language Processing*. pp. 2890-2896 (2018), https://www.aclweb.org/anthology/D18-1316

[155] Andreas, J. Good-Enough Compositional Data Augmentation. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 7556-7566 (2020), https://www.aclweb.org/anthology/2020.acl-main.676

[156] Asghar, N. Yelp dataset challenge: Review rating prediction. *ArXiv Preprint ArXiv:1605.05362*. (2016)

[157] Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M. & Weld, D. Does the Whole Exceed Its Parts? The Effect of AI Explanations on Complementary Team Performance. *Proceedings Of The 2021 CHI Conference On Human Factors In Computing Systems*. (2021), https://doi.org/10.1145/3411764.3445717

[158] Bowman, S., Angeli, G., Potts, C. & Manning, C. A large annotated corpus for learning natural language inference. *Proceedings Of The 2015 Conference On Empirical Methods In Natural Language Processing*. pp. 632-642 (2015), https://www.aclweb.org/anthology/D15-1075

[159] Ribeiro, M., Wu, T., Guestrin, C. & Singh, S. Beyond Accuracy: Behavioral Testing of NLP Models with CheckList. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 4902-4912 (2020), https://www.aclweb.org/anthology/2020.acl-main.442

[160] Chen, V., Wu, S., Ratner, A., Weng, J. & Ré, C. Slice-based Learning: A Programming Model for Residual Learning in Critical Data Slices. *Advances In Neural Information Processing Systems 32: Annual Conference On Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*. pp. 9392-9402 (2019), https://proceedings.neurips.cc/paper/2019/hash/351869bde8b9d6ad1e3090bd173f600d-Abstract.html

[161] Chomsky, N. Syntactic structures. (Walter de Gruyter,2002)

[162] Keskar, N., McCann, B., Varshney, L., Xiong, C. & Socher, R. CTRL - A Conditional Transformer Language Model for Controllable Generation. *ArXiv Preprint ArXiv:1909.05858*. (2019)

[163] Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J. & Liu, R. Plug and Play Language Models: A Simple Approach to Controlled Text Generation. *8th International Conference On Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. (2020), https://openreview.net/forum?id=H1edEyBKDS

[164] Devlin, J., Chang, M., Lee, K. & Toutanova, K. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *Proceedings Of The 2019 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long And Short Papers)*. pp. 4171-4186 (2019), https://www.aclweb.org/anthology/N19-1423

[165] Donahue, C., Lee, M. & Liang, P. Enabling Language Models to Fill in the Blanks. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 2492-2501 (2020), https://www.aclweb.org/anthology/2020.acl-main.225

[166] Dua, D., Wang, Y., Dasigi, P., Stanovsky, G., Singh, S. & Gardner, M. DROP: A Reading Comprehension Benchmark Requiring Discrete Reasoning Over Paragraphs. *Proceedings Of The 2019 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long And Short Papers)*. pp. 2368-2378 (2019), https://www.aclweb.org/anthology/N19-1246

[167] Ettinger, A. What BERT Is Not: Lessons from a New Suite of Psycholinguistic Diagnostics for Language Models. *Transactions Of The Association For Computational Linguistics*. **8** pp. 34-48 (2020), https://www.aclweb.org/anthology/2020.tacl-1.3

[168] Feder, A., Oved, N., Shalit, U. & Reichart, R. CausaLM: Causal Model Explanation Through Counterfactual Language Models. *ArXiv Preprint ArXiv:2005.13407*. (2020)

[169] Gardner, M., Artzi, Y., Basmov, V., Berant, J., Bogin, B., Chen, S., Dasigi, P., Dua, D., Elazar, Y., Gottumukkala, A., Gupta, N., Hajishirzi, H., Ilharco, G., Khashabi, D., Lin, K., Liu, J., Liu, N., Mulcaire, P., Ning, Q., Singh, S., Smith, N., Subramanian, S., Tsarfaty, R., Wallace, E., Zhang, A. & Zhou, B. Evaluating Models' Local Decision Boundaries via Contrast Sets. *Findings Of The Association For Computational Linguistics: EMNLP 2020*. pp. 1307-1323 (2020), https://www.aclweb.org/anthology/2020.findings-emnlp.117

[170] Garg, S. & Ramakrishnan, G. BAE: BERT-based Adversarial Examples for Text Classification. *Proceedings Of The 2020 Conference On Empirical Methods In Natural Language Processing (EMNLP)*. pp. 6174-6181 (2020), https://www.aclweb.org/anthology/2020.emnlp-main.498

[171] Geiger, A., Cases, I., Karttunen, L. & Potts, C. Posing Fair Generalization Tasks for Natural Language Inference. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing And The 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP)*. pp. 4485-4495 (2019), https://www.aclweb.org/anthology/D19-1456

[172] Glockner, M., Shwartz, V. & Goldberg, Y. Breaking NLI Systems with Sentences that Require Simple Lexical Inferences. *Proceedings Of The 56th Annual Meeting Of The Association For Computational Linguistics (Volume 2: Short Papers)*. pp. 650-655 (2018), https://www.aclweb.org/anthology/P18-2103

[173] Go, A., Bhayani, R. & Huang, L. Twitter sentiment classification using distant supervision. *CS224N Project Report, Stanford*. **1**, 2009 (2009)

[174] Goodwin, E., Sinha, K. & O'Donnell, T. Probing Linguistic Systematicity. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 1958-1969 (2020), https://www.aclweb.org/anthology/2020.acl-main.177

[175] Green, S., Wang, S., Chuang, J., Heer, J., Schuster, S. & Manning, C. Human Effort and Machine Learnability in Computer Aided Translation. *Proceedings Of The 2014 Conference On Empirical Methods In Natural Language Processing (EMNLP)*. pp. 1225-1236 (2014), https://www.aclweb.org/anthology/D14-1130

[176] Green, S., Chuang, J., Heer, J. & Manning, C. Predictive translation memory: A mixed-initiative system for human language translation. *Proceedings Of The 27th Annual ACM Symposium On User Interface Software And Technology*. pp. 177-187 (2014)

[177] Gururangan, S., Swayamdipta, S., Levy, O., Schwartz, R., Bowman, S. & Smith, N. Annotation Artifacts in Natural Language Inference Data. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. pp. 107-112 (2018), https://www.aclweb.org/anthology/N18-2017

[178] Hase, P. & Bansal, M. Evaluating Explainable AI: Which Algorithmic Explanations Help Users Predict Model Behavior?. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 5540-5552 (2020), https://www.aclweb.org/anthology/2020.acl-main.491

[179] Hu, Z., Yang, Z., Liang, X., Salakhutdinov, R. & Xing, E. Toward Controlled Generation of Text. *Proceedings Of The 34th International Conference On Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*. **70** pp. 1587-1596 (2017), http://proceedings.mlr.press/v70/hu17e.html

[180] Huang, W., Liu, H. & Bowman, S. Counterfactually-Augmented SNLI Training Data Does Not Yield Better Generalization Than Unaugmented Data. *Proceedings Of The First Workshop On Insights From Negative Results In NLP*. pp. 82-87 (2020), https://www.aclweb.org/anthology/2020.insights-1.13

[181] Iyyer, M., Wieting, J., Gimpel, K. & Zettlemoyer, L. Adversarial Example Generation with Syntactically Controlled Paraphrase Networks. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. pp. 1875-1885 (2018), https://www.aclweb.org/anthology/N18-1170

[182] Jain, S. & Wallace, B. Attention is not Explanation. *Proceedings Of The 2019 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long And Short Papers)*. pp. 3543-3556 (2019), https://www.aclweb.org/anthology/N19-1357

[183] Jia, R., Raghunathan, A., Göksel, K. & Liang, P. Certified Robustness to Adversarial Word Substitutions. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing And The 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP)*. pp. 4129-4142 (2019), https://www.aclweb.org/anthology/D19-1423

[184] Jiang, Y. & Bansal, M. Avoiding Reasoning Shortcuts: Adversarial Evaluation, Training, and Model Development for Multi-Hop QA. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 2726-2736 (2019), https://www.aclweb.org/anthology/P19-1262

[185] Jin, D., Jin, Z., Zhou, J. & Szolovits, P. Is BERT Really Robust? A Strong Baseline for Natural Language Attack on Text Classification and Entailment. *The Thirty-Fourth AAAI Conference On Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications Of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium On Educational Advances In Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. pp. 8018-8025 (2020), https://aaai.org/ojs/index.php/AAAI/article/view/6311

[186] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W. & Liu, P. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal Of Machine Learning Research*. **21**, 1-67 (2020), http://jmlr.org/papers/v21/20-074.html

[187] Kahneman, D. & Tversky, A. The simulation heuristic.. (Stanford Univ CA Dept of Psychology,1981)

[188] Kang, S., Jung, H., Won, D. & Lee, S. Counterfactual Explanation Based on Gradual Construction for Deep Networks. *ArXiv Preprint ArXiv:2008.01897*. (2020)

[189] Kaushik, D., Hovy, E. & Lipton, Z. Learning The Difference That Makes A Difference With Counterfactually-Augmented Data. *8th International Conference On Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. (2020), https://openreview.net/forum?id=Sklgs0NFvr

[190] Kaushik, D., Setlur, A., Hovy, E. & Lipton, Z. Explaining The Efficacy of Counterfactually-Augmented Data. *ArXiv Preprint ArXiv:2010.02114*. (2020)

[191] Kavumba, P., Inoue, N., Heinzerling, B., Singh, K., Reisert, P. & Inui, K. When Choosing Plausible Alternatives, Clever Hans can be Clever. *Proceedings Of The First Workshop On Commonsense Inference In Natural Language Processing*. pp. 33-42 (2019), https://www.aclweb.org/anthology/D19-6004

[192] Khashabi, D., Khot, T. & Sabharwal, A. More Bang for Your Buck: Natural Perturbation for Robust Question Answering. *Proceedings Of The 2020 Conference On Empirical Methods*

In Natural Language Processing (EMNLP). pp. 163-170 (2020), https://www.aclweb.org/anthology/2020.emnlp-main.12

[193] Kim, N., Patel, R., Poliak, A., Xia, P., Wang, A., McCoy, T., Tenney, I., Ross, A., Linzen, T., Van Durme, B., Bowman, S. & Pavlick, E. Probing What Different NLP Tasks Teach Machines about Function Word Comprehension. *Proceedings Of The Eighth Joint Conference On Lexical And Computational Semantics (*SEM 2019)*. pp. 235-249 (2019), https://www.aclweb.org/anthology/S19-1026

[194] Kuleshov, V., Thakoor, S., Lau, T. & Ermon, S. Adversarial examples for natural language classification problems. (2018)

[195] Kumar, V., Choudhary, A. & Cho, E. Data Augmentation using Pre-trained Transformer Models. *Proceedings Of The 2nd Workshop On Life-long Learning For Spoken Language Systems*. pp. 18-26 (2020), https://www.aclweb.org/anthology/2020.lifelongnlp-1.3

[196] Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., Toutanova, K., Jones, L., Kelcey, M., Chang, M., Dai, A., Uszkoreit, J., Le, Q. & Petrov, S. Natural Questions: A Benchmark for Question Answering Research. *Transactions Of The Association For Computational Linguistics*. **7** pp. 452-466 (2019), https://www.aclweb.org/anthology/Q19-1026

[197] Levenshtein, V. Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady*. **10** pp. 707 (1966)

[198] Li, L., Ma, R., Guo, Q., Xue, X. & Qiu, X. BERT-ATTACK: Adversarial Attack Against BERT Using BERT. *Proceedings Of The 2020 Conference On Empirical Methods In Natural Language Processing (EMNLP)*. pp. 6193-6202 (2020), https://www.aclweb.org/anthology/2020.emnlp-main.500

[199] Li, D., Zhang, Y., Peng, H., Chen, L., Brockett, C., Sun, M. & Dolan, B. Contextualized Perturbation for Textual Adversarial Attack. *Proceedings Of The 2021 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies*. pp. 5053-5069 (2021), https://www.aclweb.org/anthology/2021.naacl-main.400

[200] Li, C., Shengshuo, L., Liu, Z., Wu, X., Zhou, X. & Steinert-Threlkeld, S. Linguistically-Informed Transformations (LIT): A Method for Automatically Generating Contrast Sets. *Proceedings Of The Third BlackboxNLP Workshop On Analyzing And Interpreting Neural Networks For NLP*. pp. 126-135 (2020), https://www.aclweb.org/anthology/2020.blackboxnlp-1.12

[201] Lin, B., Zhou, W., Shen, M., Zhou, P., Bhagavatula, C., Choi, Y. & Ren, X. CommonGen: A Constrained Text Generation Challenge for Generative Commonsense Reasoning. *Findings Of The Association For Computational Linguistics: EMNLP 2020*. pp. 1823-1840 (2020), https://www.aclweb.org/anthology/2020.findings-emnlp.165

[202] Liu, N., Schwartz, R. & Smith, N. Inoculation by Fine-Tuning: A Method for Analyzing Challenge Datasets. *Proceedings Of The 2019 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long And Short Papers)*. pp. 2171-2179 (2019), https://www.aclweb.org/anthology/N19-1225

[203] Longpre, S., Wang, Y. & DuBois, C. How Effective is Task-Agnostic Data Augmentation for Pretrained Transformers?. *Findings Of The Association For Computational Linguistics: EMNLP 2020*. pp. 4401-4411 (2020), https://www.aclweb.org/anthology/2020.findings-emnlp.394

[204] Maas, A., Daly, R., Pham, P., Huang, D., Ng, A. & Potts, C. Learning Word Vectors for Sentiment Analysis. *Proceedings Of The 49th Annual Meeting Of The Association For Computational Linguistics: Human Language Technologies*. pp. 142-150 (2011), https://www.aclweb.org/anthology/P11-1015

[205] Madaan, A., Setlur, A., Parekh, T., Poczos, B., Neubig, G., Yang, Y., Salakhutdinov, R., Black, A. & Prabhumoye, S. Politeness Transfer: A Tag and Generate Approach. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 1869-1881 (2020), https://www.aclweb.org/anthology/2020.acl-main.169

[206] Reid, M. & Zhong, V. LEWIS: Levenshtein Editing for Unsupervised Text Style Transfer. *ArXiv Preprint ArXiv:2105.08206*. (2021)

[207] Malandrakis, N., Shen, M., Goyal, A., Gao, S., Sethi, A. & Metallinou, A. Controlled Text Generation for Data Augmentation in Intelligent Artificial Agents. *Proceedings Of The 3rd Workshop On Neural Generation And Translation*. pp. 90-98 (2019), https://www.aclweb.org/anthology/D19-5609

[208] Malmi, E., Severyn, A. & Rothe, S. Unsupervised Text Style Transfer with Padded Masked Language Models. *Proceedings Of The 2020 Conference On Empirical Methods In Natural Language Processing (EMNLP)*. pp. 8671-8680 (2020), https://www.aclweb.org/anthology/2020.emnlp-main.699

[209] McCoy, T., Pavlick, E. & Linzen, T. Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 3428-3448 (2019), https://www.aclweb.org/anthology/P19-1334

[210] Miller, T. Explanation in artificial intelligence: Insights from the social sciences. *Artificial Intelligence*. **267** pp. 1 - 38 (2019), http://www.sciencedirect.com/science/article/pii/S0004370218305988

[211] Miller, G. WordNet: An electronic lexical database. (MIT press,1998)

[212] Morris, J., Lifland, E., Yoo, J., Grigsby, J., Jin, D. & Qi, Y. TextAttack: A Framework for Adversarial Attacks, Data Augmentation, and Adversarial Training in NLP. *Proceedings Of The 2020 Conference On Empirical Methods In Natural Language Processing: System Demonstrations*. pp. 119-126 (2020), https://www.aclweb.org/anthology/2020.emnlp-demos.16

[213] Mozes, M., Stenetorp, P., Kleinberg, B. & Griffin, L. Frequency-Guided Word Substitutions for Detecting Textual Adversarial Examples. *Proceedings Of The 16th Conference Of The European Chapter Of The Association For Computational Linguistics: Main Volume*. pp. 171-186 (2021), https://www.aclweb.org/anthology/2021.eacl-main.13

[214] Naik, A., Ravichander, A., Sadeh, N., Rose, C. & Neubig, G. Stress Test Evaluation for Natural Language Inference. *Proceedings Of The 27th International Conference On Computational Linguistics*. pp. 2340-2353 (2018), https://www.aclweb.org/anthology/C18-1198

[215] Ni, J., Li, J. & McAuley, J. Justifying Recommendations using Distantly-Labeled Reviews and Fine-Grained Aspects. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing And The 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP)*. pp. 188-197 (2019), https://www.aclweb.org/anthology/D19-1018

[216] Nie, Y., Wang, Y. & Bansal, M. Analyzing Compositionality-Sensitivity of NLI Models. *The Thirty-Third AAAI Conference On Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications Of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium On Educational Advances In Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*. pp. 6867-6874 (2019), https://doi.org/10.1609/aaai.v33i01.33016867

[217] Lundberg, S. & Lee, S. A Unified Approach to Interpreting Model Predictions. *Advances In Neural Information Processing Systems 30: Annual Conference On Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*. pp. 4765-4774 (2017), https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html

[218] Pearl, J. Causal and Counterfactual Inference. *The Handbook Of Rationality*. pp. 1-41 (2018)

[219] Dathathri, S., Madotto, A., Lan, J., Hung, J., Frank, E., Molino, P., Yosinski, J. & Liu, R. Plug and Play Language Models: A Simple Approach to Controlled Text Generation. *8th International Conference On Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. (2020), https://openreview.net/forum?id=H1edEyBKDS

[220] Pruthi, D., Gupta, M., Dhingra, B., Neubig, G. & Lipton, Z. Learning to Deceive with Attention-Based Explanations. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 4782-4793 (2020), https://www.aclweb.org/anthology/2020.acl-main.432

[221] Qin, L., Bosselut, A., Holtzman, A., Bhagavatula, C., Clark, E. & Choi, Y. Counterfactual Story Reasoning and Generation. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing And The 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP)*. pp. 5043-5053 (2019), https://www.aclweb.org/anthology/D19-1509

[222] Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W. & Liu, P. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal Of Machine Learning Research*. **21**, 1-67 (2020), http://jmlr.org/papers/v21/20-074.html

[223] Rajpurkar, P., Zhang, J., Lopyrev, K. & Liang, P. SQuAD: 100,000+ Questions for Machine Comprehension of Text. *Proceedings Of The 2016 Conference On Empirical Methods In Natural Language Processing*. pp. 2383-2392 (2016), https://www.aclweb.org/anthology/D16-1264

[224] Rajpurkar, P., Jia, R. & Liang, P. Know What You Don't Know: Unanswerable Questions for SQuAD. *Proceedings Of The 56th Annual Meeting Of The Association For Computational Linguistics (Volume 2: Short Papers)*. pp. 784-789 (2018), https://www.aclweb.org/anthology/P18-2124

[225] Ramon, Y., Martens, D., Provost, F. & Evgeniou, T. Counterfactual explanation algorithms for behavioral and textual data. *ArXiv Preprint ArXiv:1912.01819*. (2019)

[226] Ratner, A., Bach, S., Ehrenberg, H., Fries, J., Wu, S. & Ré, C. Snorkel: Rapid training data creation with weak supervision. *Proceedings Of The VLDB Endowment*. **11**, 269-282 (2017)

[227] Ribeiro, M., Singh, S. & Guestrin, C. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. *Proceedings Of The 22nd ACM SIGKDD International Conference On Knowledge Discovery And Data Mining, San Francisco, CA, USA, August 13-17, 2016*. pp. 1135-1144 (2016), https://doi.org/10.1145/2939672.2939778

[228] Ribeiro, M., Singh, S. & Guestrin, C. Anchors: High-Precision Model-Agnostic Explanations. *Proceedings Of The Thirty-Second AAAI Conference On Artificial Intelligence, (AAAI-18), The 30th Innovative Applications Of Artificial Intelligence (IAAI-18), And The 8th AAAI Symposium On Educational Advances In Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*. pp. 1527-1535 (2018), https://www.aaai.org/ocs/index.php/AAAI/AAAI18/paper/view/16982

[229] Ribeiro, M., Singh, S. & Guestrin, C. Semantically Equivalent Adversarial Rules for Debugging NLP models. *Proceedings Of The 56th Annual Meeting Of The Association For Computational Linguistics (Volume 1: Long Papers)*. pp. 856-865 (2018), https://www.aclweb.org/anthology/P18-1079

[230] Rogers, A., Kovaleva, O. & Rumshisky, A. A Primer in BERTology: What We Know About How BERT Works. *Transactions Of The Association For Computational Linguistics*. **8** pp. 842-866 (2020), https://www.aclweb.org/anthology/2020.tacl-1.54

[231] Nakov, P., Rosenthal, S., Kozareva, Z., Stoyanov, V., Ritter, A. & Wilson, T. SemEval-2013 Task 2: Sentiment Analysis in Twitter. *Second Joint Conference On Lexical And Computational Semantics (*SEM), Volume 2: Proceedings Of The Seventh International Workshop On Semantic Evaluation (SemEval 2013)*. pp. 312-320 (2013), https://www.aclweb.org/anthology/S13-2052

[232] Sakaguchi, K., Le Bras, R., Bhagavatula, C. & Choi, Y. WinoGrande: An Adversarial Winograd Schema Challenge at Scale. *Proceedings Of The AAAI Conference On Artificial Intelligence*. **34**, 8732-8740 (2020), https://ojs.aaai.org/index.php/AAAI/article/view/6399

[233] Sanh, V., Debut, L., Chaumond, J. & Wolf, T. DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. *ArXiv*. **abs/1910.01108** (2019)

[234] Socher, R., Perelygin, A., Wu, J., Chuang, J., Manning, C., Ng, A. & Potts, C. Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank. *Proceedings Of The 2013 Conference On Empirical Methods In Natural Language Processing*. pp. 1631-1642 (2013), https://www.aclweb.org/anthology/D13-1170

[235] Song, L., Yu, X., Peng, H. & Narasimhan, K. Universal Adversarial Attacks with Natural Triggers for Text Classification. *Proceedings Of The 2021 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies*. pp. 3724-3733 (2021), https://www.aclweb.org/anthology/2021.naacl-main.291

[236] Suhr, A., Zhou, S., Zhang, A., Zhang, I., Bai, H. & Artzi, Y. A Corpus for Reasoning about Natural Language Grounded in Photographs. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 6418-6428 (2019), https://www.aclweb.org/anthology/P19-1644

[237] Tan, S., Joty, S., Kan, M. & Socher, R. It's Morphin' Time! Combating Linguistic Discrimination with Inflectional Perturbations. *Proceedings Of The 58th Annual Meeting Of The Association For Computational Linguistics*. pp. 2920-2935 (2020), https://www.aclweb.org/anthology/2020.acl-main.263

[238] Teney, D., Abbasnedjad, E. & Hengel, A. Learning What Makes a Difference from Counterfactual Examples and Gradient Supervision. *Computer Vision – ECCV 2020*. pp. 580-599 (2020)

[239] Vazirani, V. Approximation algorithms. (Springer Science Business Media,2013)

[240] Vig, J., Gehrmann, S., Belinkov, Y., Qian, S., Nevo, D., Singer, Y. & Shieber, S. Investigating Gender Bias in Language Models Using Causal Mediation Analysis. *Advances In Neural Information Processing Systems*. **33** (2020)

[241] Wang, A., Singh, A., Michael, J., Hill, F., Levy, O. & Bowman, S. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. *7th International Conference On Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. (2019), https://openreview.net/forum?id=rJ4km2R5t7

[242] Wang, Z. & Culotta, A. Robustness to Spurious Correlations in Text Classification via Automatically Generated Counterfactuals. (2020)

[243] Wei, J. & Zou, K. EDA: Easy Data Augmentation Techniques for Boosting Performance on Text Classification Tasks. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing And The 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP)*. pp. 6382-6388 (2019), https://www.aclweb.org/anthology/D19-1670

[244] Wiegreffe, S. & Pinter, Y. Attention is not not Explanation. *Proceedings Of The 2019 Conference On Empirical Methods In Natural Language Processing And The 9th International Joint Conference On Natural Language Processing (EMNLP-IJCNLP)*. pp. 11-20 (2019), https://www.aclweb.org/anthology/D19-1002

[245] Wieting, J. & Gimpel, K. ParaNMT-50M: Pushing the Limits of Paraphrastic Sentence Embeddings with Millions of Machine Translations. *Proceedings Of The 56th Annual Meeting Of The Association For Computational Linguistics (Volume 1: Long Papers)*. pp. 451-462 (2018), https://www.aclweb.org/anthology/P18-1042

[246] Williams, A., Nangia, N. & Bowman, S. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. *Proceedings Of The 2018 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. pp. 1112-1122 (2018), https://www.aclweb.org/anthology/N18-1101

[247] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q. & Rush, A. Transformers: State-of-the-Art Natural Language Processing. *Proceedings Of The 2020 Conference On Empirical Methods In Natural Language Processing: System Demonstrations*. pp. 38-45 (2020), https://www.aclweb.org/anthology/2020.emnlp-demos.6

[248] Wu, X., Lv, S., Zang, L., Han, J. & Hu, S. Conditional BERT Contextual Augmentation. *International Conference On Computational Science*. pp. 84-95 (2019)

[249] Wu, X., Lv, S., Zang, L., Han, J. & Hu, S. Conditional BERT Contextual Augmentation. *International Conference On Computational Science*. pp. 84-95 (2019)

[250] Wu, T., Ribeiro, M., Heer, J. & Weld, D. Errudite: Scalable, Reproducible, and Testable Error Analysis. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 747-763 (2019), https://www.aclweb.org/anthology/P19-1073

[251] Wu, T., Wongsuphasawat, K., Ren, D., Patel, K. & DuBois, C. Tempura: Query Analysis with Structural Templates. *CHI '20: CHI Conference On Human Factors In Computing Systems, Honolulu, HI, USA, April 25-30, 2020*. pp. 1-12 (2020), https://doi.org/10.1145/3313831.3376451

[252] Yang, Z., Hu, Z., Dyer, C., Xing, E. & Berg-Kirkpatrick, T. Unsupervised Text Style Transfer using Language Models as Discriminators. *Advances In Neural Information Processing Systems 31: Annual Conference On Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*. pp. 7298-7309 (2018), https://proceedings.neurips.cc/paper/2018/hash/398475c83b47075e8897a083e97eb9f0-Abstract.html

[253] Zhang, K. & Shasha, D. Simple fast algorithms for the editing distance between trees and related problems. *SIAM Journal On Computing*. **18**, 1245-1262 (1989)

[254] Zhang, H., Zhou, H., Miao, N. & Li, L. Generating Fluent Adversarial Examples for Natural Languages. *Proceedings Of The 57th Annual Meeting Of The Association For Computational Linguistics*. pp. 5564-5569 (2019), https://www.aclweb.org/anthology/P19-1559

[255] Zhang, Y., Baldridge, J. & He, L. PAWS: Paraphrase Adversaries from Word Scrambling. *Proceedings Of The 2019 Conference Of The North American Chapter Of The Association For Computational Linguistics: Human Language Technologies, Volume 1 (Long And Short Papers)*. pp. 1298-1308 (2019), https://www.aclweb.org/anthology/N19-1131

[256] Zhu, Y., Lu, S., Zheng, L., Guo, J., Zhang, W., Wang, J. & Yu, Y. Texygen: A Benchmarking Platform for Text Generation Models. *The 41st International ACM SIGIR Conference On Research Development In Information Retrieval, SIGIR 2018, Ann Arbor, MI, USA, July 08-12, 2018*. pp. 1097-1100 (2018), https://doi.org/10.1145/3209978.3210080

[257] Chen, Y., Li, V., Cho, K. & Bowman, S. A stable and effective learning strategy for trainable greedy decoding. *ArXiv Preprint ArXiv:1804.07915*. (2018)

[258] Gu, J., Neubig, G., Cho, K. & Li, V. Learning to translate in real-

time with neural machine translation. *ArXiv Preprint ArXiv:1610.00388*. (2016)

[259] Gu, J., Cho, K. & Li, V. Trainable greedy decoding for neural machine translation. *ArXiv Preprint ArXiv:1702.02429*. (2017)

[260] Subramani, N., Bowman, S. & Cho, K. Can unconditional language models recover arbitrary sentences?. *ArXiv Preprint ArXiv:1907.04944*. (2019)

[261] Huang, K. & Chang, K. Generating syntactically controlled paraphrases without using annotated parallel pairs. *ArXiv Preprint ArXiv:2101.10579*. (2021)