

On Zero-Shot Counterspeech Generation by LLMs

Punyajoy Saha¹, Aalok Agrawal¹, Abhik Jana², Chris Biemann³, Animesh Mukherjee¹

¹ Indian Institute of Technology, Kharagpur, ² Indian Institute of Technology, Bhubaneswar

³ Universität Hamburg, Germany

punyajoy@iitkgp.ac.in, aalokagrwal@iitkgp.ac.in, abhikjana@iitbbs.ac.in,
biemann@informatik.uni-hamburg.de, animeshm@cse.iitkgp.ac.in

Abstract

With the emergence of numerous Large Language Models (LLM), the usage of such models in various Natural Language Processing (NLP) applications is increasing extensively. Counterspeech generation is one such key task where efforts are made to develop generative models by fine-tuning LLMs with hatespeech - counterspeech pairs, but none of these attempts explores the intrinsic properties of large language models in zero-shot settings. In this work, we present a comprehensive analysis of the performances of four LLMs namely GPT-2, DialoGPT, ChatGPT and FlanT5 in zero-shot settings for counterspeech generation, which is the first of its kind. For GPT-2 and DialoGPT, we further investigate the deviation in performance with respect to the sizes (small, medium, large) of the models. On the other hand, we propose three different prompting strategies for generating different types of counterspeech and analyse the impact of such strategies on the performance of the models. Our analysis shows that there is an improvement in generation quality for two datasets (17%), however the toxicity increase (25%) with increase in model size. Considering type of model, GPT-2 and FlanT5 models are significantly better in terms of counterspeech quality but also have high toxicity as compared to DialoGPT. ChatGPT are much better at generating counter speech than other models across all metrics. In terms of prompting, we find that our proposed strategies help in improving counter speech generation across all the models.

Keywords: counterspeech, large language models, prompting

1. Introduction

Large Language Models (LLMs) like GPT-3, BARD, LLaMA are being used to produce state-of-the-art performances for numerous NLP tasks, e.g., summarisation, machine translation, text classification, etc. Despite the promising capabilities of LLMs, researchers point out their limitations in certain genres of NLP tasks like question answering (Zheng et al., 2023). Digging deep into the LLMs' behaviour for a spectrum of NLP tasks provide insights about their intrinsic properties and usability for such tasks. Therefore, in this paper, we investigate the applicability and limitations of LLMs for one of the key NLP task of *counterspeech* generation.

The rise of social media and online platforms has provided individuals with unprecedented opportunities to express themselves and engage in discussions on a global scale. Often these expressions become toxic due to bad actors spreading hate speech¹. To curtail hate speech proliferation, the moderation community has come up with the strategy of producing extensive *counterspeech* which is a direct response countering the hate/abusive speech. An example is shown below.

Hate speech: *Jews cannot be patriots, since their allegiance will always be to the state of*

¹<https://help.twitter.com/en/rules-and-policies/hateful-conduct-policy>

Israel.

Counterspeech: *You can have parents and grandparents born elsewhere and still be a patriot for the country you were born in.*

Recently, the NLP community has started exploring the usefulness of LLMs for the task of vanilla counterspeech generation as well as categorical counterspeech generation. While we find several works in the line of finetuning LLMs with hate speech - counterspeech pairs (Zhu and Bhat, 2021a), adding additional context to the LLMs (Li et al., 2022), none of them study the intrinsic properties of these models or explore prompting in a zero-shot setting to generate a specific type of counterspeech. In this paper, our contributions can be summarized as follows.

- We investigate the applicability of four LLMs (GPT-2, DialoGPT, FlanT5 and ChatGPT) for zero-shot counterspeech generation, which is the first ever attempt of this kind. We evaluate these models over four different counterspeech generation dataset – CONAN (Chung et al., 2019), CONAN-MT (Fanton et al., 2021), Reddit and Gab (Qian et al., 2019). We compare and analyse these LLMs' performances to come up with insightful observations which could be useful for the research community. We further dig into variations of a particular model in terms of size and analyse if that has

any effect in counterspeech generation.

- We propose three prompting strategies, namely *manual*, *frequency based*, and *cluster centered*. We analyze the effect of these strategies on categorical counterspeech generation.

From our detailed analysis, we make the following key observations.

Overall performance: ChatGPT outperforms all other models (DialogGPT, FlanT5, GPT-2 along with their variants) in terms of generational metrics - *gleu* (12%), *meteor* (32%) and *bleurt* (42.25%). On the other hand, counterspeech quality and argument quality improves by 120% and 35% respectively. One concerning observation is that the readability of the chatGPT generated posts reduces by 35%.

Effect of model size: With the increase of the size of the models (both for DialogGPT and GPT-2), there is an increase in *toxicity* in responses by 44%, 25%, and 30% for CONAN-MT, Reddit, and Gab respectively.

Effect of prompt type: Manual prompts perform better across *denouncing*, *facts* and *humour* type counterspeech. Cluster-centered prompts are better for *affiliation* type counterspeech for GPT2 and DialogGPT, while for the same type, frequency based prompts are better for FlanT5 and ChatGPT.

We make our code and resource used for this research public ² for reproducibility purposes.

2. Related work

Large language models: A language model estimates the probability distribution over a text. Recent advancement in the scaling of such models from a few million parameters (Merity et al., 2016) to hundred million parameters (Brown et al., 2020) and larger dataset (Gao et al., 2020a) have made them *large* language models (LLMs) performing better at a lot of downstream tasks. At this scale, the model can easily learn the downstream tasks in the few-shot as well as the zero-shot setting (Radford et al., 2019). In the zero-shot setting, we only provide template (prompts) to the models where the prompts are selected using prompt engineering.

Prompt engineering: Prompt based learning methods learn LMs that compute the probability $P(\mathbf{x}; \theta)$ of the prompt text \mathbf{x} itself and use this probability to predict the next sequence \mathbf{y} (Liu et al., 2023). So, designing the template for \mathbf{x} is an important step when these to compute the LM probabilities. Methods for text generation tasks in zero-shot setting involve manually adding prompts

to perform summarization and machine translation tasks (Radford et al., 2019). Recent methods have involved finding trigger tokens which can be used to create the prompt and do the task in zero-shot settings (Shin et al., 2020). **Counterspeech generation:** An effective strategy to mitigate hate speech is counterspeech as it does not violate the freedom of expression (Benesch et al., 2016). While the idea of countering some hateful messages is not new, the research community has recently started taking a huge interest in understanding counterspeech practices and their effectiveness while mitigating hate speech (Mathew et al., 2019). Recently, Tekiroğlu et al. proposed novel techniques to generate counterspeech using a GPT-2 model with post-facto editing by experts or annotator groups. One of the recent generation methods uses a three-stage pipeline – *Generate*, *Prune* and *Select* (GPS) to generate diverse and relevant counterspeech output (Zhu and Bhat, 2021b). Another work focuses on adding knowledge grounding to the generation pipeline (Chung et al., 2021b). A recent work further analyses several such language models and decoding strategies after finetuning them (Tekiroğlu et al., 2022). Past research has highlighted the usage of different types of counterspeech (Benesch, 2014). Two of the past works further built dataset and models for type specific counterspeech (Mathew et al., 2019; Chung et al., 2021a). Only one work focuses on generation of type specific counterspeech using generative discriminator (GED) models (Saha et al., 2022).

In all the past studies, they use a hate speech-counterspeech dataset to finetune the models first and then show their evaluation. On the other hand, our objective was to understand what these models know intrinsically. Hence, we study these LLMs in zero-shot setting to understand their intrinsic capability and compare them in terms of size and types. Second, we propose several prompting strategies which can generate type specific counterspeech as well in the zero-shot setting and further compare the different LLMs.

3. Datasets

3.1. Hatespeech-Counterspeech datasets

In order to evaluate our approach, we use four public datasets which contain hate speech and its corresponding counterspeech. The details of these datasets are noted in Table 1. Reddit and Gab datasets contain 5, 257 and 14, 614 hate speech instances, respectively (Qian et al., 2019). We use the English part of the CONAN dataset (Chung et al., 2019) which contains 408 hate speech instances. The multitarget CONAN dataset (Fantor

²https://github.com/aalokagrawal/Zeroshot_Counterspeech

et al., 2021) contains 3,718 hate speech instances. The counterspeech in Gab and Reddit datasets was written by AMT workers, whereas for CONAN, the counterspeech was written by expert NGO operators. For the CONAN-MT dataset, the pairs were generated by a generation model and later reviewed by experts.

We further made hate speech and counterspeech pairs from these datasets, such that each hate speech was associated with one counterspeech. Finally, we ended up with 3,864, 14,223, 41,580, and 5003 datapoints for CONAN, Reddit, Gab, and CONAN-MT, respectively. First, we need to divide the dataset into train and test splits. For the smaller datasets - CONAN and CONAN-MT, we find all the unique hatespeech and randomly split into train and test set with 80% for training and 20% for test set. For larger datasets - Gab and Reddit, we randomly take 500 hatespeech samples for the set of hatespeech samples. We make sure the hate speech instances in the train/validation set are not repeated in the test set to evaluate in the wild setting. For all our experiments, we use the test part of these datasets.

Dataset	Source-H	Source-C	Hate ins	# pairs	Avg. len
CONAN-MT	synthetic	expert	3,718	5,003	25.15
CONAN	synthetic	expert	408	3,864	24.07
Reddit	reddit	crowd	5,257	14,223	15.55
Gab	gab	crowd	14,614	41,580	16.05

Table 1: This table presents the source of hate speech (Source-H), source of counterspeech (Source-C), hate speech instances (Hate ins), the total pairs (# pairs) and average length of the counterspeech (Avg. len) for each of the CONAN, CONAN-MT, Reddit and Gab dataset.

3.2. Additional datasets

In this section, we describe the additional datasets required for the evaluation of the generated responses.

Counterspeech dataset: For measuring the quality of counterspeech, we use two datasets from Mathew et al. (2020a) and Chung et al. (2021a). Both these datasets have type information, while Mathew et al. (2019) in addition have non-counterspeech comments curated from YouTube. We use two variations of the counterspeech dataset.

The first variant compiles the counterspeech posts themselves. This is primarily used for **classification** of a text into a counterspeech or not. In order to align our settings with the recommendations given by Chung et al. (2021c), we place the hostile category counterspeech in the non-counterspeech part in this variant (Mathew et al., 2019). This way, we had 4,175 counterspeech comments and 9,765 non-counterspeech comments. We divide

the dataset into train:validation:test in the ratio of 8:1:1 using stratified sampling.

The second variant compiles the counterspeech types. This is primarily used for **classification** of a counterspeech into one of its types and **prompting**. We take the counterspeech posts from Mathew et al. (2019) and merge them with the counterspeech from Chung et al. (2021a) along with their types. One thing to note is that we don't utilise the posts from Mathew et al. (2019) which contain more than one strategies of counterspeech. Next, we study the definition of different types of counterspeech and select six types of counterspeech which appeared distinctive to all the authors unanimously. The statistics of the dataset finally extracted is noted in Table 2.

We divide the dataset, with 30% of it for constructing **prompts** and the rest for classification using stratified sampling. The classification dataset was further divided into train:validation:test in the ratio of 8:1:1 using stratified sampling. The prompt dataset is used to create the type prompts.

Type	Classification	Prompting	F1	clusters
hypocrisy	579	248	0.59	15
denouncing	738	316	0.85	16
humor	607	260	0.76	16
facts	1094	469	0.84	18
affiliation	163	70	0.84	16
question	227	97	0.97	18
average/total	3408	1460	0.80	16.5

Table 2: This table represents the type specific information for each of the type of counterspeech that we considered for our task. The columns *Prompting* and *Classification* represents the amount of data points used for finding prompt strategies and building classification model. *F1 score* (F1) column shows the performance of the type classifier. *Clusters* column represents the number of clusters found per type using the cluster-centered prompting strategy.

Counterargument dataset: For evaluating the counterargument quality we select a popular argument dataset (Stab et al., 2018) which has 6,317 against and 4,822 for arguments categorized into six topics. For each topic, we assume all possible pairs of arguments. From this set, we sample (without replacement) 10,000 pairs which have the same stance and 10,000 pairs which have the opposite stance. This way, we have a dataset of 60,000 argument pairs. We divide the dataset into train:validation:test in the ratio of 8:1:1 using stratified splitting.

4. Methodology

4.1. Models

We use three different model variants for understanding the zero-shot capability in counterspeech

generation.

GPT-2 (Radford et al., 2019) is trained on a large dataset called WebText. The dataset contains the textual content found in 45 million links shared by users on Reddit (Trinh and Le, 2018). Note that, WebText is not directly sourced from Reddit itself, but rather consists of data derived from outbound links posted on Reddit. This language model is trained with the objective of predicting the next word, given all the previous words within some text. The model aims at maximising $p(x) = \prod_{i=1}^n p(x_i|x, \dots, x_{i-1})$, for our experiments we use all three different versions of GPT-2 – 117M (small), 345M (medium), and 762M (large) parameters from this link³.

DialoGPT (Zhang et al., 2020) is trained on a large corpus consisting of English Reddit dialogues. The corpus consists of 147 million instances of dialogues, collected over a period of 12 years. Unlike GPT-2, this model generates better dialogue-like responses to any given prompt. In this model, along with ground truth response $T = x_1, \dots, x_n$, we also have a dialogue utterance history S . The model aims at maximising $p(T|S) = p(x_1|S) \prod_{i=2}^n p(x_i|S, \dots, x_{i-1})$. For our experiment, we use all three different versions of DialoGPT - 117M (small), 345M (medium), and 762M (large) parameters from this link⁴.

FlanT5 (Chung et al., 2022) is a T5 (Raffel et al., 2020) model that has been finetuned on a multi-task mixture of supervised tasks and for which each task is converted into a text-to-text format. They use instruction finetuning (Ouyang et al., 2022) procedure on each of these tasks, as well as chain-of-thought (CoT) prompting (Wei et al., 2022). Overall, the authors show that using such a framework improves the results across various benchmarks over the T5 versions.

ChatGPT (OpenAI, 2022) is trained with a GPT 3.5 model using Reinforcement Learning from Human Feedback (RLHF), using a similar method as InstructGPT (Ouyang et al., 2022) but with slight differences in the data collection setup. ChatGPT performs exceptionally well in question-answering scenarios. Beyond its capability of being a conversational tool, many attempts have been made to evaluate the quality of ChatGPT-generated texts in various domains (Yang et al., 2023).

4.2. Prompting strategies

Type specific generation is another challenge in counterspeech generation (Mathew et al., 2019). In this regard, the first few words of the counterspeech can be essential for the type we want to generate. We propose three different prompting

strategies to generate the first few words (prompts). These prompts help us in controlling the type of the counterspeech generated by the above LLMs.

Manual prompting: In this strategy, two authors experienced in hate speech detection research read through the prompts dataset and created 2-3 possible beginnings (prompts) for each type of counterspeech. These prompts guide the model to generate the appropriate type of counterspeech. For manual prompts, we do not set a hard limit on the number of words but asked the authors to make them small.

Frequency based prompting: In this strategy, we collect the beginning four words as a sub-string from each counterspeech of each type and cluster using exact matching of the sub-strings. These sub-strings represent the prompts for a particular type. We take the top five prompts based on their frequency for each type.

Cluster centered prompting: In this strategy, we first pass all the counterspeech from the prompting dataset through sentence embedding model `all-mpnet-base-v2`⁵. For each type of counterspeech we then cluster the embeddings using K -means clustering. We decide the number of cluster for each type of counterspeech using the elbow method (Marutho et al., 2018). We note the number of clusters for each type in 2. The average clusters are ~ 16 per type. Out of these clusters we select the top 10 clusters using their cluster size. Then we select top three sentences per cluster which lie closest to the cluster center. These sentences act as a representative of that cluster. The final prompts per cluster are comprising the beginning 4 words as a sub-string from each of these three sentences. This way we collect 30 prompts per type in total.

We note one instance of prompts for each strategy per type in the Table 3.

Type	M	CC	FB
F	This is a fact	the myth that muslims	The vast majority of
Hy	In contradiction	i am wondering have	If you are really
Hu	This is funny	i bet she got	Must be hard for
A	I also belong	i am jewish and	I am a christian
Q	Are you aware of	how do you know	Why do you think
D	Please do not say	why is the hate	Why is this a

Table 3: One instance of prompt using each prompt strategy per type. Each column represents the different prompt strategies and each row represents a particular type of counterspeech. M: Manual, CC: Cluster centered, FB: Frequency based. F: Facts, Hy: Hypocrisy, Hu: Humour, Af: Affiliation, Q: Questions, D: Denouncing.

³https://huggingface.co/docs/transformers/model_doc/gpt2 ⁵<https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

⁴<https://github.com/microsoft/DialoGPT>

4.3. Experimental setup

Our experimental setup comprises generation of a post when we pass the hate speech as an input to the LLMs. In case of type prompts, we further add the `type-prompt` at the end of the hate speech as input. For ChatGPT API, we use a similar setting where we just pass the hate speech without any prompts to the ChatGPT API; for the type prompts we add “start the counterspeech with following `type-prompt`”. We extract the `type-prompt` by randomly selecting a prompt from the given list of prompts for a particular type.

Sequences are generated up to a minimum length of 40 and a maximum of 60 tokens. We set the generation length a bit higher as compared to average length of the counterspeech (as noted in Table 1) in order to allow the model more freedom to generate. In addition, for DialoGPT, FlanT5 and GPT-2 models, top k sampling and top p sampling (aka nucleus sampling) are used while generating from our trained models. At each generation step, all the generated tokens are ranked according to their probabilities, and the top 100 most probable tokens are selected for broad distributions. In the case of narrow distributions, all the tokens are included until their CDF is 0.92 following the recommendations by Holtzman et al. (2020). The temperature is 1.2, and the repetition penalty is 3.5. For ChatGPT API, we add a system message “you are a helpful assistant that generates counterspeech”. We also keep top p sampling, repetition penalty and the temperature same.

5. Evaluation metrics

Generation metrics: To measure the generation quality, we use different standard metrics. We use `gleu` (Wu et al., 2016) and `meteor` (Banerjee and Lavie, 2005) to measure how similar the generated counterspeech are to the ground truth references. We also measure if the LLMs generates diverse and novel counterspeech. For this purpose, we use the `diversity` and `novelty` metrics from existing literature (Wang and Wan, 2018). In addition, we also report one of the recent generation metrics, `bleurt` (Sellam et al., 2020). Note that, we do not use the BLEU (Papineni et al., 2002) score because it has some undesirable properties when used for single sentences, as it is designed to be a corpus-specific measure (Wu et al., 2016). Further, the reader might notice negative scores in the case of `bleurt` metric. This is not unnatural since the `bleurt`, unlike BLEU, is not calibrated. For more information, refer here⁶.

Engagement prediction metrics: We use the DialogRPT model (Gao et al., 2020b) to predict the

human feedback of the counterspeech generated using the following metrics – *width*: the number of direct replies to the given reply, *depth*: the maximum length of dialog after this turn, and *updown*: the number of upvotes minus the number of downvotes.

This metric can help us in identifying how engaging the generated counterspeech is, which is another important characteristic, as noted by Benesch et al. (2016). To calculate the engagement metric, we pass the `hate speech-counterspeech` pair to the model, which provides a score between 0 and 1 representing the engagement in terms of upvotes/width/height. This will denote the engagement probability of that metric for the given counterspeech.

Quality measurement metrics: We deploy various third-party classifiers to evaluate the quality of the generated responses. To calculate the scores, we pass the generated counterspeech through the model and get the logit scores, which are passed through a softmax layer. The metrics used for evaluation are listed below.

- *Argument:* In order to evaluate the argument characteristic of the generated response, we use a `roberta-base-uncased` model⁷ fine-tuned on the argument dataset (Stab et al., 2018). Given this model, we pass each generated response through the classifier to predict a confidence score, which would denote the argument quality.
- *Counterargument:* In order to evaluate the counterargument characteristic of the generated response, we use a `bert-base-uncased` model trained on the counterargument dataset defined in section 3.2. We achieve an F1-score of 0.62 on the test set of this dataset. Given this model, we pass each of the hate speech and the generated response through the classifier to predict a confidence score, which would denote the counterargument quality.
- *Counterspeech:* In order to evaluate the counterspeech quality of the generated responses, we use a `bert-base-uncased` model trained on the counterspeech dataset introduced in section 3.2. We achieve an F1-score of 0.7 on the test set of this dataset. Given this model, we pass each generated response through the classifier to predict a confidence score, which denotes the quality of the counterspeech.
- *Toxicity:* We use the HateXplain model (Mathew et al., 2020b) trained on two classes – toxic and non-toxic⁸. We report

⁷<https://huggingface.co/chk1a/roberta-argument>

⁸<https://huggingface.co/Hate-speech-CNERG/bert-base-uncased-hatexplain-rationale-two>

⁶<https://github.com/google-research/bleurt/issues/1>

model	gleu	met	div	nov	blrt	cs	c_arg	arg	tox (↓)	fre
CONAN_MT										
DGPT-(s)	0.07	0.08	0.84	0.84	-1.13	0.15	0.54	0.26	0.24	65.21
DGPT-(m)	0.07	0.08	0.84	0.84	-1.16	0.16	0.54	0.22	0.19	72.46
DGPT-(l)	0.07	0.09	0.83	0.83	-1.14	0.15	0.59	0.23	0.28	66.07
GPT-2	0.06	0.11	0.82	0.84	-1.02	0.56	0.50	0.38	0.13	43.24
GPT-2-(m)	0.06	0.10	0.83	0.85	-1.05	0.51	0.49	0.36	0.18	44.25
GPT-2-(l)	0.06	0.11	0.83	0.84	-1.04	0.48	0.47	0.36	0.19	43.27
flan-T5-(s)	0.08	0.13	0.81	0.82	-0.94	0.40	0.56	0.48	0.18	61.21
flan-T5-(b)	0.08	0.13	0.81	0.81	-0.90	0.43	0.49	0.47	0.21	60.65
flan-T5-(l)	0.08	0.12	0.82	0.82	-0.96	0.41	0.46	0.43	0.18	58.08
ChatGPT	0.09	0.17	0.66	0.80	-0.53	0.95	0.64	0.51	0.15	29.89
CONAN										
DGPT-(s)	0.09	0.11	0.88	0.87	-1.21	0.15	0.58	0.20	0.31	60.67
DGPT-(m)	0.09	0.11	0.88	0.87	-1.23	0.09	0.54	0.15	0.24	70.95
DGPT-(l)	0.09	0.12	0.86	0.86	-1.20	0.08	0.65	0.21	0.37	63.57
GPT-2	0.08	0.15	0.85	0.86	-1.06	0.48	0.55	0.37	0.20	41.09
GPT-2-(m)	0.08	0.15	0.85	0.86	-1.06	0.34	0.54	0.38	0.23	44.65
GPT-2-(l)	0.08	0.15	0.85	0.86	-1.08	0.43	0.51	0.35	0.21	43.73
Flan-T5-(s)	0.10	0.17	0.84	0.84	-1.86	0.33	0.56	0.40	0.26	61.59
Flan-T5-(b)	0.10	0.17	0.84	0.84	-1.84	0.33	0.52	0.44	0.25	53.48
Flan-T5-(l)	0.10	0.17	0.84	0.84	-0.98	0.31	0.58	0.42	0.22	55.32
ChatGPT	0.12	0.23	0.69	0.81	-0.63	0.89	0.64	0.44	0.23	32.05
Gab										
DGPT-(s)	0.05	0.07	0.87	0.86	-1.26	0.06	0.53	0.06	0.09	58.18
DGPT-(m)	0.05	0.07	0.86	0.86	-1.26	0.08	0.55	0.06	0.09	57.00
DGPT-(l)	0.05	0.09	0.85	0.84	-1.28	0.07	0.56	0.06	0.09	59.25
GPT-2	0.05	0.12	0.83	0.85	-1.37	0.31	0.53	0.19	0.15	58.16
GPT-2-(m)	0.05	0.12	0.84	0.85	-1.37	0.28	0.54	0.19	0.19	58.47
GPT-2-(l)	0.05	0.12	0.83	0.85	-1.36	0.28	0.53	0.19	0.19	55.44
FlanT5-(s)	0.06	0.11	0.84	0.84	-1.37	0.24	0.56	0.22	0.16	67.10
FlanT5-(b)	0.06	0.11	0.84	0.83	-1.35	0.23	0.50	0.21	0.20	68.33
FlanT5-(l)	0.06	0.11	0.84	0.83	-1.34	0.26	0.52	0.19	0.16	63.79
ChatGPT	0.08	0.17	0.64	0.80	-0.71	0.90	0.64	0.26	0.12	29.77
Reddit										
DGPT-(s)	0.05	0.06	0.87	0.88	-1.22	0.07	0.59	0.07	0.07	30.52
DGPT-(m)	0.05	0.07	0.87	0.87	-1.21	0.08	0.55	0.06	0.07	58.24
DGPT-(l)	0.06	0.08	0.86	0.86	-1.25	0.08	0.61	0.07	0.07	62.26
GPT-2	0.05	0.12	0.82	0.86	-1.34	0.36	0.57	0.21	0.12	55.06
GPT-2-(m)	0.05	0.12	0.83	0.86	-1.35	0.35	0.56	0.22	0.14	52.91
GPT-2-(l)	0.05	0.12	0.83	0.86	-1.34	0.35	0.55	0.21	0.16	52.88
FlanT5-(s)	0.06	0.12	0.83	0.84	-1.35	0.31	0.57	0.26	0.12	73.82
FlanT5-(b)	0.06	0.11	0.84	0.84	-1.34	0.29	0.51	0.22	0.16	70.51
FlanT5-(l)	0.06	0.11	0.84	0.84	-1.32	0.34	0.53	0.20	0.11	70.99
ChatGPT	0.08	0.17	0.67	0.81	-0.77	0.85	0.50	0.26	0.13	29.12

Table 4: Evaluation of responses generated by each model for each counterspeech generation dataset in terms of generation, engagement and quality metrics. The first column denotes which model is being used for zero-shot evaluation. DialoGPT (DGPT) and GPT-2 has s, m and l suffixes which represent 117M, 345M and 762M parameter sizes, and FlanT5 has s, b and l suffixes which represent 80M, 250M and 750M parameter sizes. For evaluating generation we measure the average gleu, meteor (met), bleurt (blrt), novelty (nov) and diversity (div). Engagement metrics consist of upvote, width, and depth. For quality, we utilise the counterspeech (cs), argument (arg), counter argument (c_arg) and toxicity (tox) scores and readability scores (fre). **Bold** denotes the best scores and higher scores denote better performance except for toxicity.

the confidence between $[0, 1]$ for the toxic class. This metric is important because a toxic counterspeech might escalate the discussion.

Readability: We further evaluate the readability of the counter speech generated. We use a popular readability metric - Fleish Reading Ease (Farr et al., 1951) (fre). It gives a score between 0-100. **Type classifier:** In order to evaluate the type specific generation, we train a bert-base-uncased model on the type based counterspeech data points mentioned in section 3.2 using a multi-class classification strategy. Overall, we achieve an average macro F1-score 0.80. Among the types, we achieve a macro F1-score ~ 0.80 for denouncing, humor, facts and affiliation. Hypocrisy is hardest to clas-

sify with an F1-score of 0.59 and questions are the easiest to classify with an F1-score of 0.97.

6. Results

6.1. Vanilla generation

Here we discuss the evaluation results for the zero-shot evaluation of various models for the vanilla generation setting.

Does counterspeech generation depend on model size in zero-shot setting? We compare the small, medium (base) and large sizes of three different variations of models, i.e., DialoGPT, FlanT5 and GPT-2. We note the percentage change be-

tween the largest and smallest model. We observe that for the synthetic datasets, i.e., CONAN and CONAN-MT the change is not significant in terms of generation metrics. In terms of counterspeech quality, we see a drop of 42% for the DialoGPT model for the CONAN dataset, whereas there is a drop of 13% for GPT-2 in case of CONAN-MT dataset. In terms of counterargument quality, we notice a drop of 6-9% across all the models except DialoGPT for CONAN dataset. Surprisingly, for CONAN-MT dataset, the toxicity increases by 44% for the GPT-2 models as we increase the model size. For the real world datasets, i.e., Gab and Reddit, there is a significant increase in toxicity (25-30%) as we increase the size of GPT-2 model. On the other hand, we find that there is an improvement in the generation quality for the DialoGPT in terms of `gleu` and `meteor` metrics. In addition, we find the readability of the counterspeech generated by DialoGPT on the Reddit dataset increases dramatically (100%) with increase in the size of the model. The size trend for Flan-T5 is not consistent although it performs quite poorly compared to DialoGPT and GPT-2.

Does counterspeech generation depend on model type in zero-shot setting? We compare DialoGPT, GPT-2 and FlanT5 models since they vary in their architecture, pretraining (finetuning) strategies and dataset used for pre-training. In terms of synthetic datasets, FlanT5 models are better in terms of `gleu` (30%) and `meteor` (50%) and GPT-2 models are better in `bleurt` (8-15%). In terms of other metrics, we find that counterspeech quality is far better for GPT-2 models (200%) than other models. DialoGPT models are also better in terms of readability, but at the same time are more toxic. For the real world datasets, i.e., Gab and Reddit, DialoGPT models are better in terms of `bleurt` (7-10%), GPT-2 models are better in `meteor` and FlanT5 models are better in terms of `gleu` (20%). In terms of other metrics, we find that counterspeech quality and argument improves by 2–3 times for GPT-2 and FlanT5 models, but they are also higher in terms of toxicity.

6.2. Type specific generation

In this part, we evaluate the type specific generation using type prompts. We run the counterspeech classifier (described in section 5) over the post generated for a particular type and measure the ratio of posts which the classifier classifies as the same type. We name this metric as `type precision`. The type prompts and their extraction procedure is mentioned in section 4.2. Since we do not observe much change in the type specific generation for small, medium and large versions of the DialoGPT, FlanT5 GPT-2, we present their average performance in the Table 6. For *affiliation*, GPT-2

Model	Hate speech	Counterspeech
FlanT5	Muslims have no place in civilised Europe. They should go home.	We shall not discriminate on the grounds of race, religion, colour, culture or national origin. Our society defends and protects its own dignity.
DGPT	We are being invaded by Muslims. They are taking over our countries	lol s I hope we all get this reference.The irony is lost at the end of that sentence isn't it...
ChatGPT	If you are a Jew you are for sure involved in 9/11 planning. This is all about Race and 'Enlightenment'. You are all guilty until proven innocent	It is unfair and untrue to suggest that all Jews were involved in the planning of 9/11...

Table 5: Counterspeech generation by different models in a zero-shot setting for vanilla generation. The posts are taken from one of the available datasets. For all models except ChatGPT, we use the `1` version of the models.

and DialoGPT perform better with cluster centered prompts across all the datasets, while FlanT5 performs better with frequency based prompts. Cluster centered prompts improve the baseline type precision by 0.41 and 0.32 units for GPT-2 and DialoGPT respectively and frequency based prompts improve the FlanT5 by 0.16 units. For *denouncing*, GPT-2, DialoGPT and FlanT5 perform better with manual prompts. The type precision of FlanT5 increases by 0.16 units while the other models increases by 0.36 units. *Facts* type⁹ are one of the categories where the GPT-2 and FlanT5 performs quite better than DialoGPT type precision. Overall, the baseline FlanT5 model is also good for generating fact type statements for CONAN and CONAN-MT datasets, with scoring ~ 0.7 in type precision. Here, also manual prompts are the best, and they improve the baseline type precision by 0.20, 0.15 and 0.13 units for GPT-2, DialoGPT and FlanT5 respectively. For *humour*, we find again that the manual prompts are the best. They improve the baseline type precision by 0.185, 0.27 and 0.12 units for GPT-2, DialoGPT and FlanT5 respectively. For *hypocrisy* and *questions*, there is no one prompting strategy that works for both the models and all the datasets. Cluster centered prompts are better for GPT-2 (0.10 units) and manual prompts are better for DialoGPT (0.15 units). Surprisingly, none of the models reach 0.1 type precision for the *question* type. We have added some examples of type specific generation

⁹We don't claim that these models are generating factual statements, but fact-styled statements.

model	prompt	aff	den	fac	hum	hyp	qu
CONAN_MT							
GPT-2	base	0.04	0.04	0.60	0.03	0.29	0.00
	manual	0.06	0.40	0.78	0.20	0.21	0.01
	freq	0.06	0.07	0.77	0.10	0.21	0.05
DGPT	base	0.04	0.10	0.29	0.19	0.39	0.00
	manual	0.10	0.45	0.46	0.46	0.51	0.01
	freq	0.10	0.15	0.44	0.34	0.50	0.08
FlanT5	base	0.04	0.06	0.73	0.03	0.14	0.00
	manual	0.15	0.23	0.81	0.15	0.10	0.00
	freq	0.26	0.06	0.80	0.07	0.21	0.00
ChatGPT	base	0.02	0.22	0.75	0.00	0.01	0.00
	manual	0.03	0.40	0.81	0.00	0.02	0.00
	freq	0.51	0.31	0.94	0.00	0.06	0.01
GPT-2	base	0.02	0.04	0.65	0.01	0.27	0.00
	manual	0.05	0.37	0.81	0.13	0.19	0.01
	freq	0.04	0.07	0.79	0.05	0.18	0.04
DGPT	base	0.02	0.09	0.28	0.20	0.41	0.00
	manual	0.07	0.44	0.42	0.46	0.52	0.01
	freq	0.07	0.16	0.42	0.37	0.50	0.07
FlanT5	base	0.02	0.08	0.75	0.02	0.12	0.00
	manual	0.10	0.23	0.79	0.14	0.12	0.00
	freq	0.18	0.07	0.80	0.05	0.21	0.00
ChatGPT	base	0.04	0.64	0.28	0.00	0.10	0.00
	manual	0.02	0.39	0.93	0.00	0.02	0.00
	freq	0.52	0.30	0.96	0.00	0.06	0.00
GPT-2	base	0.08	0.07	0.27	0.17	0.41	0.00
	manual	0.20	0.48	0.51	0.40	0.48	0.01
	freq	0.20	0.16	0.50	0.26	0.50	0.08
DGPT	base	0.03	0.07	0.14	0.47	0.30	0.00
	manual	0.08	0.43	0.30	0.75	0.52	0.01
	freq	0.08	0.14	0.29	0.64	0.52	0.07
FlanT5	base	0.09	0.12	0.26	0.23	0.30	0.00
	manual	0.17	0.27	0.44	0.34	0.30	0.00
	freq	0.23	0.12	0.41	0.25	0.30	0.01
ChatGPT	base	0.07	0.39	0.44	0.00	0.10	0.00
	manual	0.07	0.60	0.75	0.02	0.10	0.00
	freq	0.54	0.48	0.75	0.00	0.13	0.01
GPT-2	base	0.08	0.09	0.26	0.18	0.40	0.00
	manual	0.23	0.49	0.51	0.40	0.47	0.01
	freq	0.23	0.19	0.48	0.28	0.46	0.07
DGPT	base	0.03	0.08	0.12	0.47	0.29	0.00
	manual	0.09	0.46	0.27	0.73	0.52	0.01
	freq	0.09	0.16	0.27	0.64	0.53	0.07
FlanT5	base	0.10	0.12	0.24	0.25	0.29	0.00
	manual	0.17	0.28	0.43	0.37	0.29	0.00
	freq	0.23	0.11	0.41	0.28	0.29	0.00
ChatGPT	base	0.02	0.09	0.88	0.00	0.01	0.00
	manual	0.07	0.71	0.62	0.01	0.07	0.00
	freq	0.52	0.66	0.65	0.00	0.09	0.00
GPT-2	base	0.08	0.09	0.26	0.18	0.40	0.00
	manual	0.23	0.49	0.51	0.40	0.47	0.01
	freq	0.23	0.19	0.48	0.28	0.46	0.07
DGPT	base	0.03	0.08	0.12	0.47	0.29	0.00
	manual	0.09	0.46	0.27	0.73	0.52	0.01
	freq	0.09	0.16	0.27	0.64	0.53	0.07
FlanT5	base	0.10	0.12	0.24	0.25	0.29	0.00
	manual	0.17	0.28	0.43	0.37	0.29	0.00
	freq	0.23	0.11	0.41	0.28	0.29	0.00
ChatGPT	base	0.02	0.09	0.88	0.00	0.01	0.00
	manual	0.07	0.71	0.62	0.01	0.07	0.00
	freq	0.52	0.66	0.65	0.00	0.09	0.00

Table 6: Evaluation of responses generated by each model and prompt strategy. The first column denotes which model is being used for zero-shot evaluation. DialoGPT, GPT-2 and FlanT5 are averaged across three parameter sizes. The second column denotes the prompt strategy out of manual_prompt (manual), frequency based (freq), cluster centered (cluster) being used, where baseline (base) represents no prompt strategy. The next six columns represent the type-precision for the counterspeech generated by each model + type_prompt. aff: affiliation, den: denouncing, fac: facts, hum: humour, hyp: hypocrisy, qu: question. **Bold** denotes the best scores and higher scores denote better performance.

Model (Type)	Hate speech	Counterspeech
GPT-2 (aff)	Blacks are just an inferior bloodline to our own. We should remove them.	i am muslim but my black brothers don't respect us as true Muslims, and there is another cause of that.. so how do I make their hearts be happy then they can accept me? the one thing we have in common with other races...is faith - you know all over this world...every people has a belief"
GPT-2 (den)	No, only the Severely Retarded get the job.	i think maybe prejudice against people with that disease does drive a lot of discrimination in their lives - they've got some serious issues and it's very hard to help them
DialoGPT (hum)	Never said it was, faggot	This is funny. uhh you know who loves the attention and self hate speech as well? Donald Trump,
DialoGPT (hyp)	People of a different race, including transgenders, are very few. With some effort we can contain them.	In contradiction to the people living in fear and ignorance. gender neutrality exists even though those who want it aren't happy about their pronouns?
FlanT5 (aff)	We should expel Muslims, everyone of them out of the country.	I have no issue with Muslims. I just disagree that they are the bad guys and there are lots of other countries like my country in which they won't go, so we shouldn't do anything about them.
ChatGPT (facts)	Women should not be allowed to vote. They are subhuman.	I am a Muslim and in our faith, both men and women are valued equally. Discriminating against one gender is unjustified, as it goes against the principles of justice that Islam teaches us to follow ..

Table 7: Counterspeech generated by different models in a zero-shot setting for type specific generation where the type is denoted in the brackets. For each type, we select the prompt strategy based on the performance of that model on that type. aff: affiliation, den: denouncing, fac: facts, hum: humour, hyp: hypocrisy, qu: question. For all models except ChatGPT, we use their larger versions.

in the Table 7.

6.3. Generation using ChatGPT

In this part, we look at the counterspeech generated by ChatGPT. In terms of vanilla generation, as noted in Table 4, we notice that ChatGPT performs better than other models in terms of generation quality. It improves the *gleu* (12%), *meteor* (32%) and *bleurt* (42.25%). Among other metrics, ChatGPT improves the counterspeech quality by 120% and improves the argument quality by 27%. The toxicity scores are comparable, although they are slightly higher than the best models. One interesting point is that readability of the ChatGPT texts reduces by $\sim 35\%$. We also note some counterspeech generated by ChatGPT and compare them with the other models in Table 5. In terms of type specific generation as noted in Table 6, we find that for *affiliation*, frequency based prompts improves the baseline type precision by 0.49 units across all datasets. Other than that, other types do not have a consistent best prompt strategy that works across all the categories. Interestingly, for some cases like denouncing type for CONAN, the model performance worsens if we introduce any prompt. Further ChatGPT model performs best in fact-type counterspeech scoring close to 0.9 in three out of four datasets. ChatGPT struggles for the types—humour, hypocrisy, and questions even in presence of type prompts and their type precisions rarely reach above 0.1 across datasets.

7. Discussion and conclusion

In this work, we presented a thorough analysis of the performance of LLMs in a zero-shot setting to generate counterspeech. This is to understand what these models are capable of intrinsically without training on hate speech-counterspeech pairs. We explored the models namely DialoGPT, GPT-2 and FlanT5 in terms of their size and further extended the experiments to ChatGPT.

In case of vanilla experiments, we find these models show some promise in generating counterspeech in zero-shot setting, with ChatGPT outperforming the other two models. Further, we do not see many changes in the variations among the GPT-2 models. The improvement for ChatGPT is also visible as we manually evaluated the generation (shown in Table 5). Hence, like other tasks, the emergent behaviour is only visible when we increase the scale (i.e., GPT-2 \rightarrow ChatGPT).

Next, we proposed three prompting strategies to generate categorical counterspeech and analysed the applicability of all these different models for the same. We find that carefully designed manual prompts are better than our proposed automatic methods. Although, these prompts are able to control the generations for a smaller model like

DialoGPT, it fails for ChatGPT except for some specific types like facts. This opens the door for future research, which can focus on two directions (a) design of better prompting strategies, and (b) improve models like ChatGPT to better capitalise on these prompts to benefit the type specific generations for LLMs.

Do the metrics correlate with human judgements? While we present most of our results with automatic metrics, it is important to understand if they correlate with human judgements. We took one referential (*bleurt*) and non referential metric (counterspeech). For each metric, 25 samples were extracted from each tail of the predicted metric values. We present these to expert researchers in the hate speech domain and ask them to rate the quality of counterspeech from 1-5, 5 being the best and 1 being the worst. For the counterspeech metric (nominal ratings, ordinal human evaluations), the Point-biserial correlation coefficient (Linacre and Rasch, 2008) was 0.45. For the *bleurt* metric (continuous ratings, ordinal human evaluations), the Spearman's rank correlation was 0.56. These results highlight the consistency between automated metrics and human judgments, affirming their reliability.

Ethics statement

Hate speech is a complex phenomenon. While the language generation methods are better than before, it is still very far from generating coherent and meaningful counterspeech (Bender et al., 2021). Further, they are very unreliable. There are multiple cases where the chatbots turned hateful when deployed without supervision, leading to their shutting down¹⁰. Hence, we advocate against the deployment of fully automatic pipelines for countering hate speech (de los Riscos and D'Haro, 2021). Based on the current progress in this pipeline, active participation of the counter speakers is required to generate relevant counterspeech. Our efforts to study the counterspeech generation by these automatic models can help in further improvement in the counterspeech generation pipeline and better inform the counter speakers.

8. Bibliographical References

Satanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and*

¹⁰<https://www.cbsnews.com/news/microsoft-shuts-down-ai-chatbot-after-it-turned-into-racist-nazi/>

- extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Susan Benesch. 2014. Countering dangerous speech: New ideas for genocide prevention. *Washington, DC: United States Holocaust Memorial Museum*.
- Susan Benesch, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Lucas Wright. 2016. Considerations for successful counter-speech. *A report for Public Safety Canada under the Kanishka Project*. Accessed November, 25:2020.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*.
- Yi-Ling Chung, Marco Guerini, and Rodrigo Agerri. 2021a. Multilingual counter narrative type classification. *arXiv preprint arXiv:2109.13664*.
- Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroglu, and Marco Guerini. 2019. Conan-counter narratives through nichesourcing: a multilingual dataset of responses to fight online hate speech. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2829.
- Yi-Ling Chung, Serra Sinem Tekiroğlu, and Marco Guerini. 2021b. [Towards knowledge-grounded counter narrative generation for hate speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 899–914, Online. Association for Computational Linguistics.
- Yi-Ling Chung, Serra Sinem Tekiroğlu, and Marco Guerini. 2021c. [Towards knowledge-grounded counter narrative generation for hate speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 899–914, Online. Association for Computational Linguistics.
- Online. Association for Computational Linguistics.
- Agustín Manuel de los Riscos and Luis Fernando D’Haro. 2021. [ToxicBot: A Conversational Agent to Fight Online Hate Speech](#), chapter Conversational Dialogue Systems for the Next Decade. Springer Singapore.
- Margherita Fanton, Helena Bonaldi, Serra Sinem Tekiroğlu, and Marco Guerini. 2021. [Human-in-the-loop for data collection: a multi-target counter narrative dataset to fight online hate speech](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3226–3240, Online. Association for Computational Linguistics.
- James N Farr, James J Jenkins, and Donald G Paterson. 1951. Simplification of flesch reading ease formula. *Journal of applied psychology*, 35(5):333.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. 2020a. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*.
- Xiang Gao, Yizhe Zhang, Michel Galley, Chris Brockett, and Bill Dolan. 2020b. [Dialogue response ranking training with large-scale human feedback data](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 386–395, Online. Association for Computational Linguistics.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *International Conference on Learning Representations*.
- Yu Li, Baolin Peng, Yelong Shen, Yi Mao, Lars Liden, Zhou Yu, and Jianfeng Gao. 2022. [Knowledge-grounded dialogue generation with a unified knowledge representation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 206–218, Seattle, United States. Association for Computational Linguistics.
- JM Linacre and G Rasch. 2008. The expected value of a point-biserial (or similar) correlation. *Rasch Measurement Transactions*, 22(1):1154.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig.

2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9):1–35.
- Dhendra Marutho, Sunarna Hendra Handaka, Ekaprana Wijaya, and Muljono. 2018. [The determination of cluster number at k-mean using elbow method and purity evaluation on headline news](#). In *2018 International Seminar on Application for Technology of Information and Communication*, pages 533–538.
- Binny Mathew, Navish Kumar, Pawan Goyal, and Animesh Mukherjee. 2020a. [Interaction dynamics between hate and counter users on twitter](#). In *Proceedings of the 7th ACM IKDD CoDS and 25th COMAD, CoDS COMAD 2020*, page 116–124, New York, NY, USA. Association for Computing Machinery.
- Binny Mathew, Punyajoy Saha, Hardik Tharad, Subham Rajgaria, Prajwal Singhania, Suman Kalyan Maity, Pawan Goyal, and Animesh Mukherjee. 2019. Thou shalt not hate: Countering online hate speech. In *Proceedings of the international AAI conference on web and social media*, volume 13, pages 369–380.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2020b. Hatexplain: A benchmark dataset for explainable hate speech detection. *arXiv preprint arXiv:2012.10289*.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. In *International Conference on Learning Representations*.
- OpenAI. 2022. Introducing chatgpt. <https://openai.com/blog/chatgpt>. (Accessed on 06/11/2023).
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Jing Qian, Anna Bethke, Yinyin Liu, Elizabeth Belding, and William Yang Wang. 2019. A benchmark dataset for learning to intervene in online hate speech. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4755–4764.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Punyajoy Saha, Kanishk Singh, Adarsh Kumar, Binny Mathew, and Animesh Mukherjee. 2022. [Countergerdi: A controllable approach to generate polite, detoxified and emotional counter-speech](#). In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, pages 5157–5163. International Joint Conferences on Artificial Intelligence Organization. AI for Good.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. Bleurt: Learning robust metrics for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online. Association for Computational Linguistics.
- Christian Stab, Tristan Miller, Benjamin Schiller, Pranav Rai, and Iryna Gurevych. 2018. [Cross-topic argument mining from heterogeneous sources](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3664–3674, Brussels, Belgium. Association for Computational Linguistics.
- Serra Sinem Tekiroğlu, Helena Bonaldi, Margherita Fanton, and Marco Guerini. 2022. Using pre-trained language models for producing counter narratives against hate speech: a comparative study. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3099–3114.
- Serra Sinem Tekiroğlu, Yi-Ling Chung, and Marco Guerini. 2020. Generating counter narratives against online hate speech: Data and strategies.

- In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1177–1190.
- Trieu H Trinh and Quoc V Le. 2018. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*.
- Ke Wang and Xiaojun Wan. 2018. Sentigan: Generating sentimental texts via mixture adversarial networks. In *IJCAI*, pages 4446–4452.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Bing Yin, and Xia Hu. 2023. Harnessing the power of llms in practice: A survey on chatgpt and beyond. *arXiv preprint arXiv:2304.13712*.
- Yizhe Zhang, Siqi Sun, Michel Galley, Yen-Chun Chen, Chris Brockett, Xiang Gao, Jianfeng Gao, Jingjing Liu, and William B Dolan. 2020. Dialogpt: Large-scale generative pre-training for conversational response generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 270–278.
- Shen Zheng, Jie Huang, and Kevin Chen-Chuan Chang. 2023. Why does chatgpt fall short in answering questions faithfully? *arXiv preprint arXiv:2304.10513*.
- Wanzheng Zhu and Suma Bhat. 2021a. [Generate, prune, select: A pipeline for counterspeech generation against online hate speech](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 134–149, Online. Association for Computational Linguistics.
- Wanzheng Zhu and Suma Bhat. 2021b. Generate, prune, select: A pipeline for counterspeech generation against online hate speech. *arXiv preprint arXiv:2106.01625*.